Experiments in a Software Aided Multiagent System

Shu-Heng Chen and Chung-Ching Tai

AI-ECON Research Center
Department of Economics
National Chengchi University
Taipei, Taiwan
chchen@nccu.edu.tw
elliot@aiecon.org

Abstract. Agent-based Methodology (ABM) is becoming indispensable for the inter-disciplinary study of social and economic complex adaptive systems. The essence of ABM lies in the notion of autonomous agents whose behavior may evolve endogenously and can generate and mimic the corresponding complex system dynamics that the ABM is studying. Over the past decade, many Computational Intelligence (CI) methods have been applied to the design of autonomous agents, in particular, their adaptive scheme. This design issue is non-trivial since the chosen adaptive schemes usually have great impact on the generated system dynamics. Robert Lucas, one of the most influential modern economic theorists, has suggested using laboratories with human agents, also known as Experimental Economics, to help solving the selection issue. While this is a promising approach, laboratories used in the current experimental economics is not computationally equipped to meet the demands of the selection task. This paper attempts to materialize Lucas' suggestion by establishing a laboratory where human subjects are equipped with the computational power that satisfies the *computational equivalence* conditions.

1 Introduction

The use of agent-based simulation to study complex adaptive systems has become increasingly popular. Its significance has been well demonstrated by a series of recent conferences and journals exclusively devoted to this subject. One of the major issues in agent-based simulation is *agent engineering*. In the agents literature, many have reported that the simulation results are highly dependent upon how agents are designed and how agents learn and adapt, i.e. *agent engineering does matter*. Hence, many work have been devoted to the research of agent engineering.

One important topic that agent engineering addresses is the *robustness* of agent-based simulation results. This is normally done by evaluating whether different learning algorithms used in the simulation would lead to different implications. While the famous *KISS principle*[1] has provided a simple approach, e.g. *reinforcement learning*, for agent engineering, other sophisticated approaches, e.g., *genetic algorithms*, are also very popular in the agents literature. In fact, many techniques within the field of *computational intelligence* have been used for agent engineering. Given such a wide variety of techniques, there is a great need for a guideline on how to select an appropriate technique for agent engineering.

Recently, there is a guideline for techniques selection based on empirical observations. This guideline gives two possibilities. The first one uses data from field studies, surveys or census, while the second one uses data from laboratories with human subjects. In economics, laboratories data are becoming increasingly available from the study of *Experimental Economics* and *Behavioral Economics*. Since data from field studies, surveys or census are becoming difficult to obtain, it is inevitable for economists to use the second type of data to conduct research. This approach is known as *Lucas criterion* ([21, 5]). Lucas's criterion suggests that a comparison of the behavior of adaptive schemes with behavior observed in laboratory experiments involving human subjects can facilitate the choice of a particular adaptive scheme.

There are already many studies that grounded their agent engineering in the spirit of *Lucas's criterion*. For example, [1] used two versions of genetic algorithm (*basic GA* and *augmented GA*) to implement agent engineering [1]. She reported that the simulation results from basic GA give individual quantities and prices exhibited fluctuations over the entire duration and did not result in convergence to the rational expected equilibrium values, which was inconsistent with the experimental results involving human subjects. In contrast, the results of the augmented GA showed convergence to the rational expected equilibrium values, and were able to capture several features of the experimental behavior of human subjects better than other simple learning algorithms. According to Lucas criterion, the augmented GA was justified as an appropriate adaptive scheme. For more application of the Lucas criterion to justify the use of the GA or some of its specific versions, please refer to [2] and [3].

In addition to Lucas criterion, empirical evidence from experiments has been used to examine if reinforcement learning adequately describes the way people behave. For example, [13] gives an evaluation of reinforcement learning versus some of its competitive alternatives, such as *direction learning* and *belief learning*, in experimental asymmetric-information games.

The purpose of this paper is to answer the following question:

To what extent experimental economics or behavioral economics can help build agent engineering? Does the proposed guideline really have a solid foundation?

Our hypothesis is that experimental economics and behavioral economics have their limits and can not solve the foundation issue of agent engineering. The published guideline therefore requires revision. We shall use scientific method to prove the hypothesis in this research.

---

[1] The KISS principle was first proposed in [4]. It stands for "*Keep it simple, stupid*."

2 Computational Equivalence

Our hypothesis is based on the rationale that different learning algorithm requires different *computational resources* and that difference, to the best of our knowledge, has not been addressed in the experimental design including human subjects. As a result, the claim that an adaptive scheme (generated by one type of learning algorithm) is superior over others (generated by other types of learning algorithm), based on their simulation results, is not always valid. One scenario can be that the adaptive schemes become unavailable to human subjects due to the lack of high-performance computing facilities required by the adaptive schemes. Another scenario is for on-line transaction experiments where a very limited amount of time is allocated to the agents (human subjects) to process the data before making a decision. Such constraint makes it impossible for the human agents to carry out computationally-intensive adaptive schemes, such as *genetic programming* or *fuzzy neural networks*. Without considering the issue of computational resources, any comparative study based on experimental results is meaningless.

Let's take the *double-auction experiment* as an example. In this auction, both sellers and buyers submit bids which are then ranked from the highest to lowest to generate demand and supply profiles. Based on the profiles, the maximum quantity exchanged can be determined by matching selling offers (starting with the lowest price and moving up) with demand bids (starting with the highest price and moving down). Experiments based on double auction have a long history in experimental economics ([29]). While the experiments involving human subjects started 40 years ago, serious work on agent engineering is only just beginning. The zero-intelligence (ZI) trader was introduced by [15]. It was then extended into the zero-intelligence plus (ZIP) trader by [10]. More sophisticated agents based on human-written programs were considered by [26] and [27], which also motivated [7] to use genetic programming to make traders autonomous. Another related study is that by [12] who used the genetic algorithm to evolve traders. Given such a variety, can we answer which one provides better agent engineering in light of the experiments conducted by [29]? Can the experimental economics of the 1970s and 1980s resolve the selection issue?

Clearly not. ZIP, GA and GP differ in their required computational resources. Consequently, as mentioned above, in a computationally-poor environment, it is not surprising to see that the behavior of human subjects is quite similar to the agent-based simulation using simple heuristics, such as the *ZIP Plus* scheme, and may be quite distinctive from the one using genetic programming. However, that result alone cannot effectively lend support to the superiority of *ZIP Plus* over GP. Without knowing this critical relation, experiments involving human subjects would simply be too arbitrary to be a foundation of agent engineering. Unfortunately, this subtle point has not received sufficient attention on the part of either experimental economists or agent-based computational economists.

Experimental economics was developed in an age where decision supports provided by intensive computation was not available, whereas agent-based computational economics was cultivated in an era accompanied by increasing effi-

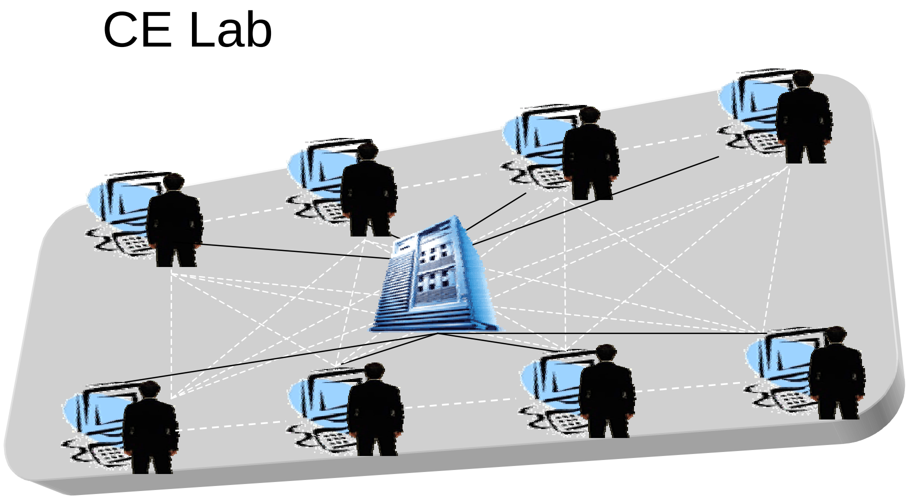

Fig. 1. A Web of the CE Lab
In the CE laboratory, each human participant will be connected to the "society" through a personal computer (client) and the server. The interactive dynamics of all participants then takes place through the web.

ciency in terms of both software and hardware. This sharp difference in computational background does not make the use of the former as the foundation of the latter as obvious as one might have thought initially. To address the pivotal questions, such as whether genetic algorithms (or any other CI tools) represent an essential learning process on the past of humans, the laboratory involving human subjects must be upgraded to such a degree that GA (or any other CI tool) can be effectively executed for human subjects. We shall call this condition *computational equivalence*. In other words, *unless the condition of computational equivalence is satisfied, experimental economics cannot serve as a way of defining a principle of agent engineering*. It is misleading to claim the superiority of one adaptive scheme over another simply by citing the experimental results observed in the laboratory which are not computationally equivalent.

Given this discussion, we propose a computer laboratory which is built in line with the condition of computational equivalence. In this laboratory, all human subjects can *choose* to follow any adaptive scheme (CI tool) which is already installed in the lab and is made available to the end users. Of course, they can also choose their own preferable ways of making decisions, e.g., relying on their simple heuristics, if they fail to see the benefits of using the sophisticated adaptive schemes. The whole idea is depicted in Figures 1 to 3.

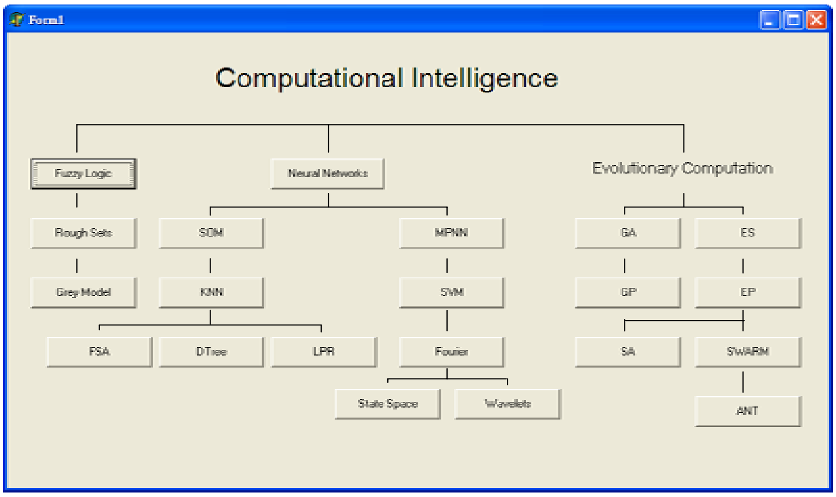

Fig. 2. Menu of Adaptive Schemes (CI Tools)
CI tools are made available for each human agent. From the computer screen, each human agent can choose his/her favorite adaptive scheme to learn from and adapt to the surrounding environment. The screen is presented to them in a very friendly manner. In this case, different classes of CI tools are grouped in an organization chart.

The whole laboratory is webbed by a network as shown in Figure 1. This web provides the basic platform for running experiments involving human subjects. Each human subject (real agent) is connected to at least one computer (*client*). The participation of the real agents in the experiments, e.g., through submitting a limit order in a stock market experiment ([30]), proceeds via the client to the *server* and further to the whole web. Figure 2 is an expanding picture of what the human subject can see from his/her computer screen. Via this client machine, the real agent can acquire all the information pertinent to the experiment. The real agent may not have direct eye contact with other real agents since these clients are not necessarily located in the same room. Nonetheless, in the case where these clients are distributed in different places, real agents may still have eye contact with other participants via the attached camera.[2]

The idea of *computational equivalence* is exemplified in Figures 2 and 3. In addition to the basic information providing by the experiment, the client also provides real agents with a computationally-rich environment such that the

[2] To the best of our knowledge, experiments relying on eye contact are very limited. One of the examples is the famous *ultimatum game*.

agents are able to perform some non-trivial computation before they finalize their decision. Here, real agents can make their choice of the adaptation scheme that they would like to follow. For example, in the context of the double auction experiment, they can follow the GA to submit a bid or ask. If they *all* make such a choice, then the market dynamics will be pretty much similar to what [12] predicted. They can also follow the GP to learn a trading strategy first, and based on that strategy submit a bid or ask. Again, if they *all* make such a choice, then the market dynamics will be close to what [7] predicted. Of course, there is no reason why they should all choose the same adaptive scheme, as each can make a choice independently from the menu shown in Figure 2.

One recent research trend in computational intelligence involves the extensive use of hybrid systems, in which many tools work synergetically together as a *multi-agent system*. The CE lab also allows participants to organize their own hybrid systems from a set of CI tools. For example, a genetic neural fuzzy system ([24]) is shown in Figure 3.

Now, we can define what *computational equivalence* means. For an agent-based computational model, if each of the human agents behaves by *choosing exactly the same* form of agent engineering as the model suggests, then the experimental results, either individual or aggregate, will be the same as the agent-based simulation results, at least in a statistical sense. In other words, in a way quite different from the current directional relationship between experimental economics and agent-based computational economics, we are not only interested in replicating experimental results based on agent-based simulations, but what we have also said here is that agent-based simulation results can also be replicated by experiments if the "right" form of agent engineering is chosen by the human subjects in the experiment. This is the essence of *computational equivalence*.

The contributions of computational equivalence to complex adaptive economic systems are two-fold. *First*, it provides a *competing explanation* for the failure of some adaptive schemes to replicate the experimental results. In the past, such failures would usually be largely attributed to the irrelevancy of the chosen adaptive schemes. However, a competing explanation nowadays is that the laboratory is not well designed to satisfy the CE condition. Had it been so, those adaptive schemes might have proved their relevance in describing how humans behave. *Second*, through the CE lab, we can see that the agent-based computational economic model is no longer just a theoretical representation of the reality, but it can also serve as a blueprint for an economic system to be realized in the future. For example, [18] is not just an agent-based simulation of the fish market. If the fish market somehow becomes automated one day in the future and all market participants can get access to a computationally-rich environment to form their decisions, then [18] can be considered to be a blueprint of the real fish market.

Let us elaborate on the first point above using the example of *genetic programming*, one of the very popular CI tools. The following is a list of criticisms

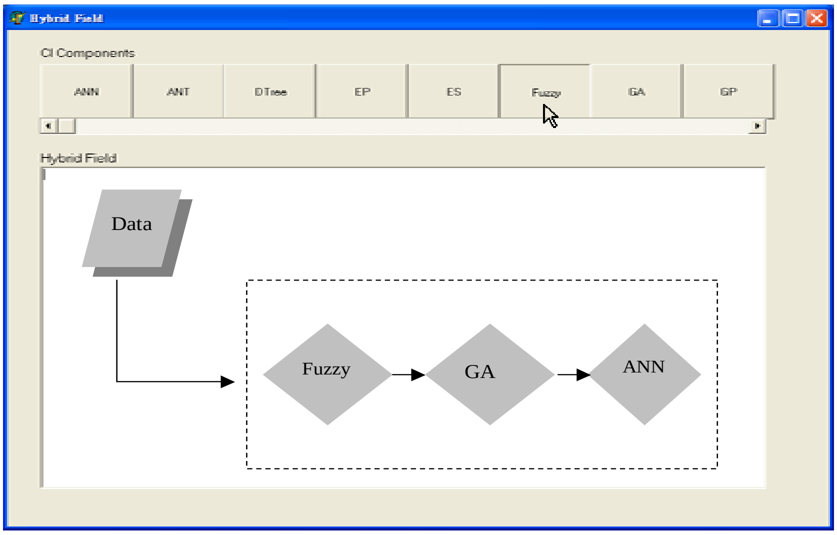

Fig. 3. Advanced Menu of Adaptive Schemes (CI Tools)
A human agent may combine several different CI tools into a hybrid system to learn and adapt. For example, a genetic neural fuzzy system is shown here.

voiced by the journal referees on several of the authors' submissions. For the convenience of discussion, they are numbered as follows:

1. One can see why genetic programming (an extension of genetic algorithms) might work for computer problem-solving. However, genetic programming is not well grounded in considerations of human behavior.
2. There just isn't sufficient justification in the paper to support the view that what we have is a model for a population of agents learning over time (if they seriously wanted to push this view they would, for instance, need to tell readers how to interpret the cross-over operation).
3. The regular GP representation and presence of operators often lead to overly complex and unreliable solutions. Such solutions, comprised of complex combinations of functions and indicators, have often been difficult to understand and interpret. Deciphering the winning programs might be an impossible task in many applications. It might be impossible to use these rules in order to understand the process by which humans behave. The usefulness of the GP method for fundamental research thus seems quite restricted *a priori*.

As we shall argue here, none of these arguments is true when computational equivalence is imposed. Objection (1) is the most general type of criticism. It is true that, up to the present, genetic programming is not well grounded in considerations of human behavior, but the main reason for this is that we do not have a lab which can make it easier for human subjects to use GP in forming decisions. As was discussed before, GP is a highly computationally-demanding CI tool. How can we expect human subjects to seriously base their decisions on GP without first reducing their computational load to a reasonable degree? So, clearly, Objection (1) cannot be valid because it violates the computational equivalence condition.

Objection (2) is also similar to Objection (1), but is more technical-oriented. This kind of comment is frequently made in relation to other CI tools. In general, this means that a CI tool cannot be regarded as a sensible human learning and adaptation process if its major operator, e.g., the *cross-over operator* in GP, is empirically not analogous to human behavior. Without acknowledging computational equivalence, this argument is also misplaced. Whether or not a CI tool can be an effective description of human behavior has nothing to do with whether the society has an innate sense of that operator. On the other hand, managers who use GP to make predictions may not need to be convinced that its technical operators have empirical social meaning. As long as human behavior is concerned, the key issue is whether they would *believe* that GP can enhance their quality decisions. This is an empirical question, and can only be solved with a lab based on computational equivalence. A full understanding of its technical details may be neither a necessary nor a sufficient condition for the formation of this belief.

Just as in the case of Objection (2), Objection (3) proposes another possible reason why people may not use GP. This criticism is commonly shared in relation to other CI tools characterized by complex and nonlinear behavior. Our response to Objection (3) is similar to that to Objection (2). Whether or not a comprehensive understanding is crucial for humans to accept a state-of-the-art technology can only be answered empirically. It is quite common in behavioral economics for human choices to not necessarily be rational or scientific. Various biases have been well established empirically in behavioral economics.

In summary, the relevance of computational intelligence to human adaptive behavior cannot be appropriately studied without a CE laboratory (a lab that satisfies the computational equivalence condition). The first step of this research is, therefore, to build a CE lab, and then start to run subsequent experiments.

3 Complex Systems Comprising Human and Software Agents

The CE lab described above not only serves as a starting point toward a foundation of agent engineering but also function as a platform to study the behavior of *electronic markets*. In this type of markets, all decisions are either made by humans aided by software or are automated completely. The former is exactly the architecture depicted in Figure 4, whereas the latter is equivalent to a CE lab where all human agents are replaced by software agents as shown in the same figure.

There are already many economic and financial applications of these two architecture. For example, in Internet auctions, software agents, such as *esnipe* and *auctionblitz*, support human bidders with routine task to improve performance precision ([23, 25]). Another example is in financial markets where programmed trades on the computerized trading platform can be used to help day traders in determining their bid and ask decisions ([32]). Our CE lab can be viewed as a simple model of these complex adaptive systems, and, via agent-based simulation, can be used to enrich our understanding of this type of systems.

However, our CE Lab is not designed only for specific applications, such as electronic markets. Its framework is very general so that we can use it to address a very fundamental question: *What will happen when the participated human agents are equipped with CI tools complex adaptive system?* We try to answer this question using a picture. Figure 4 gives four different types of agent-based models. On the left of the top row, we have an agent-based model that is completely composed of human agents. On the right of the top row, we have a model that is purely composed of software agents. Two models on the bottom are agent-based models that are composed of both human and software agents. The difference between the two is that human agents on the left model are aware of the existence of software agent while the human agents on the right model are not.

Most experimental economic models only consider human agents, whereas most agent-based computational models only have software agents. While the original motivation of agent-based models has been to understand the equivalent systems composed of human agents, human participation and interface with the system and with the artificial agents have been excluded. In other words, human agents have been completely replaced by artificial agents.[3] Recently, the literature has started to look at the issue of interaction between software agents and human agents in complex adaptive systems, such as auction games ([11, 17, 16, 28, 33]), oligopoly games ([22]) and the stock market ([19]). Nonetheless, in these setups, human agents are only exposed to limited computational resources, and are not equipped with high-performance computing; therefore, the potential feedback relation between human agents and the system, including the artificial agents, has not been greatly exploited. We believe that the CE lab is the first step toward a formal study of the fundamental question posed above.

Our CE lab enables us to conduct experiments based on the four different types of agent-based models and hence makes it possible for us to perform a comparative study on the complex dynamics of the four agent-based models. Moreover, in addition to the comparative study, the CE lab also allows us to examine the possible *structural changes* or *regime changes* when a society that

[3] Having said that, we would like to point out that the modern definition of artificial intelligence has already given upon the dream of replicating human behavior. Now, a more realistic and also interesting definition is based on the *team work* cooperatively performed by artificial agents and human agents ([20]).

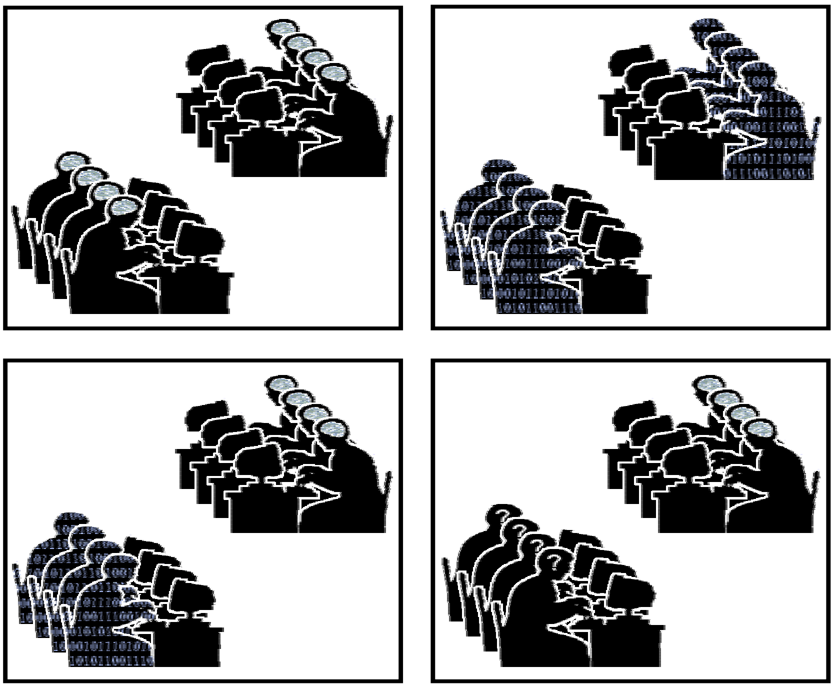

Fig. 4. Complex Systems Composed of Human and/or Software Agents

is solely composed of human agents is gradually or suddenly transforming itself into a society with a mixed population of software agents and human agents. The main points of interest are as follows:

– First, in general, how do humans react to the nonlinear complex dynamics? Would they tend to behave in a more complex manner when situated in such a complex environment?
– Second, if they do behave in a more complex manner[4], what are the consequences of such behavior? Will the aggregate dynamics become more (or less) complex and nonlinear?
– Third, what are the additional impacts of the inclusion of artificial agents (human agents) on an already-existing society composed of only human agents (artificial agents)? Do artificial agents on markets influence human behavior? ([16])

[4] For example, they incorporate some of the CI tools into their decisions.

Regarding question (1), an interesting hypothesis is that humans may not necessarily respond to a complex environment with complex behavior. On the contrary, they may just follow a *rule of thumb*. This *simplicity hypothesis* is partially supported by some theoretical and empirical work.[5] For example, the winning strategy in trading tournament involving many agents is simple ([26, 27]). In trading, the simple trading strategy *buy-and-hold* was found to exhibit superior performance to many other sophisticated strategies. In making predictions, the simple predictor *random walk* was also shown to predict more accurately than other sophisticated predictors. Nonetheless, these findings are not strong enough to discourage people to devote a great deal of effort to develop sophisticated nonlinear predictors as well as trading strategies. Therefore, what determines people's search intensity for a "best" perception (representation) of complex nonlinear phenomena remains an interesting issue.

Through these experimental designs and the associated agent-based simulations of the CE lab, we can watch and study how human agents react to a complex nonlinear environment. In particular, it helps us to understand under what circumstances human agents tend to behave in ways that are more complex than simple.[6] More generally, can we find an effective characterization or implicit constraint of agents' choices of adaptive schemes? For example, we need to ask whether they tend to prefer qualitative schemes over quantitative schemes, linguistic schemes over crisp schemes, and simple (comprehensible) but sub-optimal schemes over complex (incomprehensible) but accurate schemes,..., etc. Needless to say, these are important questions that need to be addressed before we can lay a solid foundation for agent engineering.

One important recent development in agent-based social simulation has occurred in the use of *natural language* ([17, 31]). People frequently and routinely use natural language or linguistic values, such as high, low, and so on, to describe their perception, demands, expectations, and decisions. Some psychologists have argued that our ability to process information efficiently is the outcome of applying *fuzzy logic* as part of our thought process. Evidence on human reasoning and human thought processes supports the hypothesis that at least some categories of human thought are definitely fuzzy. Yet, early agent-based economic models have assumed that agents' adaptive behavior is *crisp*. [31] made progress in this direction by using the *genetic-fuzzy classifier system* (GFCS) to model traders' adaptive behavior in an artificial stock market. [31] provided a good illustration of the *non-equivalence* between the acknowledgement of the *cognitive constraint* and the assumption of *simple agents*. It is well-known that the human mind is notoriously bad at intuitively comprehending exponential growth. However, there is no evidence that traders on Wall Street are simple-minded. [31] recog-

[5] This hypothesis can be loosely connected to the famous *Occam's razor* or the *parsimony principle*.

[6] While [8] simulates the evolving complexity of software agents in the artificial stock markets, the empirical counterpart of human behavior is not available in the literature due to the lack of an appropriate lab, such as the CE Lab.

nized the difference, and appropriately applied the GFCS to *lessen* the agents' reasoning load via the use of natural language.

The thing that concerns the second issue is the casual relation between the complexity of macro dynamics and the complexity of micro- dynamics, or simply the *emergent phenomena*. While a series of studies regarding the emergent phenomena were conducted in the past in the ACE context ([14, 9, 6]), there is also no empirical evidence in relation to experimental economics.

A particularly interesting thing is that the micro behavior can sometimes be quite different from the macro behavior. Both the work done by [14] on the cobweb model and by [9] on the asset pricing model has shown that the time series of the market price (an aggregate variable) follows a simple stochastic process. However, there is no simple description of the population dynamics of individual behavior. The simple stochastic price behavior was, in effect, generated by a great diversity of agents whose behavior was constantly changing. [9] proposed a measure for the *complexity* of an agent's behavior and a measure of the *diversity* of an agent's complexity, and it was found that both measures can vary quite widely, regardless of the simple aggregate price behavior.

In addition, using micro-structure data, [6] initiated an approach to study the *emergent property*. By that definition, they found that a series of aggregate properties, such as the efficient market hypothesis, the rational expectations hypothesis, the price-volume relation and the sunspot effect, which were proved by rigorous econometrics tests, were generated by a majority of agents who did not believe in these properties. Once again, our understanding of the micro behavior does not lead to a consistent prediction of the macro behavior. The latter is simply not just the linear scaling-up of the former. Conventional economics tends to defend the policy issues concerned with the individual's welfare, e.g., the national annuity program, based on the macroeconometrics tests, e.g., the permanent income hypothesis. Agent-based macro-economics may invalidate this approach due to emergent properties.

By using the CE lab, one can vary the computational resources under different experiments. For example, in one experiment, human agents are exposed only to simple adaptive schemes, whereas in the other they can gain access to sophisticated adaptive schemes. In this way, we can examine whether a society of naive agents will tend to result in less complex aggregate dynamics than a society of sophisticated agents.

The third issue in a sense concerns the effects of the appearance of software agents on human agents. When human agents are explicitly informed of the presence of anonymous software agents in the system, would they behave differently as compared to the case where such a presence is uncertain ([16])? Furthermore, when human agents are provided with more detailed information regarding how software agents behave, such as their adaptive schemes, would that affect their own choice of adaptive schemes? This question is particularly relevant for electronic trading systems, such as *ebay*. In a sense, this question can be viewed in terms of the *socio-psychological impact* on human behavior in the presence of interacting machine intelligence. An equally important issue is concerned with the associated market dynamics and efficiency ([7]).

4 Concluding Remarks

This paper has two aims. The first is to build a laboratory with human subjects, which satisfies the *computational equivalence* condition. The second is so that the lab, referred as the *CE lab*, can then serve as a platform to integrate the current research in experimental economics, behavioral economics (finance) and agent-based computational economics. These three fields share a common research goal with regard to the relevance and significance of adaptive behavior to economic dynamics. While, from the agent-based computational economics, it is now quite clear that aggregate dynamics can crucially depend on the learning dynamics (the so-called *agent engineering* has not been effectively resolved. In particular, it cannot be resolved solely on the basis of experimental economics in the way Robert Lucas suggested. This is because most labs in which experiments involving human subjects are conducted do not provide subjects with high-performance computing facilities. Consequently, many adaptive schemes studied in agent-based computational economics are virtually impossible for human agents to compute, and the empirical relevance of adaptive schemes such as fuzzy logic, neural networks, and genetic algorithms is beyond the current research capacity of experimental economics.

Computational equivalence is about *replictability*. It basically requires that the lab have the same computational power as ACE generally has. To achieve this goal, software agents are introduced to and work with human agents in the lab, as suggested by the MIT approach (to artificial intelligence). By means of computational equivalence, what is done by the software agents (autonomous agents) in ACE may in principle be replicated by human agents in this lab. Only when replictability is guaranteed, can one ground the foundation of agent engineering in experiments involving human agents.

Leaving aside ACE, the failure to incorporate these computational intelligence techniques into the current experimental economics naturally limits the computational complexity of human decision rules. This limitation can be a real concern when the human behavior which experimental economics tries to study is largely inspired by a computationally-rich environment. Specifically, in an era of electronic commerce, when more and more automated trading techniques are being made available to human agents, questions regarding the dynamics and the efficiency of different auction designs may no longer be properly answered by the experiments conducted in the conventional lab. However, the use of the CE lab would help. Furthermore, more intriguing questions may arise when human agents do not only interact (play, compete) with human agents, but also with the possible presence of software agents, as is often seen in the case of eBay or Nasdaq. What are the impacts of software agents on human agents? What is the effect on market dynamics and efficiency when autonomous agents are introduced to the markets? More generally, how do human agents adapt to their digital surroundings and what are the consequential dynamics? The CE lab, by effectively integrating ACE and experimental economics, provides us with a starting place to explore the richness of the nonlinear complex digital economy.

References

1. Arifovic, J. (1994), "Genetic Algorithm Learning and the Cobweb Model," *Journal of Economic Dynamics and Control* **18(1)**, pp. 3–28.
2. Arifovic, J. (1995), "Genetic Algorithms and Inflationary Economies," *Journal of Monetary Economics* **36(1)**, pp. 219–243.
3. Arifovic, J., and M. Maschek (2003), "Expectations and Currency Crisis–An Experimental Approach," paper presented at the *9th International Conference on Computing in Economics and Finance*, University of Washington, Seattle, July 11–13.
4. Axelrod, R. (1997), "Advancing the Art of Simulation in the Social Sciences," in R. Conte, R. Hegselmann, and P. Terna (Eds.), *Simulating Social Phenomena*, Lecture Notes in Economic and Mathematical Systems, pp. 21–40.
5. Chen, S.-H., and C.-H. Yeh (1996), "Genetic Programming Learning and the Cobweb Model," in P. Angeline (Ed.), *Advances in Genetic Programming*, Vol. 2, Chap. 22. MIT Press, Cambridge, MA, 443–466.
6. Chen, S.-H. and C.-C. Liao (2004), "Agent-Based Computational Modeling of the Stock Price-Volume Relation," forthcoming in *Information Sciences*.
7. Chen, S.-H. and C.-C. Tai (2003), "Trading Restrictions, Price Dynamics and Allocative Efficiency in Double Auction Markets: Analysis Based on Agent-Based Modeling and Simulations," forthcoming in *Advances in Complex Systems*.
8. Chen, S.-H., and C.-H. Yeh (2001), "Evolving Traders and the Business School with Genetic Programming: A New Architecture of the Agent-Based Artificial Stock Market," *Journal of Economic Dynamics and Control* **25**, pp. 363–393.
9. Chen, S.-H., and C.-H. Yeh (2002), "On the Emergent Properties of Artificial Stock Markets," in *Journal of Economic Behavior and Organization*, Vol. 49, pp. 217–129.
10. Cliff, D. (1997), "Minimal-Intelligence Agents for Bargaining Behaviors in Market-Based Environments," *HP Technical Report*, HPL-97-91, 1997.
11. Das, R., J. Hanson, J. Kephart, and G. Tesauro (2001), "Agent-Human Interactions in the Continuous Double Auction," *Proceedings of the International Joint Conference on Artificial Intelligence*.
12. Dawid, H. (1999), "On the Convergence of Genetic Learning in a Double Auction Market," *Journal of Economic Dynamics and Control*, 23, pp. 1545–1567.
13. Feltovich, N. (2000), "Reinforcement-based vs. Belief-based Learning Models in Experimental Asymmetric-Information Games," *Econometrica*, Vol. 68, pp. 605–641.
14. Franke, R. (1998), "Coevolution and Stable Adjustments in the Cobweb Model," *Journal of Evolutionary Economics* **8(4)**, pp. 383–406.
15. Gode, D. K., and S. Sunders (1993), "Allocative Efficiency of Market with Zero-Intelligence Trader: Market as a Partial Substitute for Individual Rationality," *Journal of Political Economy*, Vol. 101, No. 1, pp. 119–137.
16. Grossklags, J. and C. Schmidt (2003), "Interaction of Human Traders and Artificial Agents on Double Auction Markets: Simulations and Laboratory Experiments," in K. Chen, et al. (Eds.), *Proceedings of 7th Information Sciences*, pp. 1269–1272.
17. He, M., H.-F. Leung, and N. Jennings (2002), "A Fuzzy Logic-Based Bidding Strategy for Autonomous Agents in Continuous Double Auctions," *IEEE Transactions on Knowledge and Data Engineering*.

18. Kirman, A. P., and N. Vriend (2001), "Evolving Market Structure: An ACE Model of Price Dispersion and Loyalty," *Journal of Economic Dynamics and Control* **25(3-4)**, pp. 459–502.
19. Kurumatani, K., T. Yamamoto, H. Kawamura, and A. Ohuchi (2004), "Market Micro-Structure Analysis by Multi-Agent Simulation in an X-EconomyX Comparison among Technical Indices," *Information Sciences*, forthcoming.
20. Lieberman, H. (1996), "Software Agents: The MIT Approach," Invited speech delivered at the *7th European Workshop on Modelling Autonomous Agents in a Multi-Agent World* (MAAMAW'96), Eindhoven, The Netherlands, Jan 22–25, 1996.
21. Lucas, R. E., Jr. (1986), "Adaptive Behavior and Economic Theory," *Journal of Business* **59**, pp. 401–426.
22. Midgley, D., R. Marks, and L. Cooper (1997), "Breeding Competitive Strategies," *Management Science* **43 (3)**, pp. 257–275.
23. Ockenfels, A. and A. Roth (2002), "The Timing of Bids in Internet Auctions: Market Design, Bidder Behavior, and Artificial Agents," *Artificial Intelligence Magazine*, Fall, pp. 79–87.
24. Ringhut, E. and S. Kooths (2003), "Modeling Expectations with GENEFER - An Artificial Intelligence Approach," *Computational Economics*, Vol. **21**, pp. 173–294.
25. Roth, A., and A. Ockenfels (2002), "Last Minute Bidding and the Rules of Ending Second Price Auctions: Evidence from eBay and Amazon Auctions on the Internet," *American Economic Review* **92(4)**, pp. 1093–1103.
26. Rust, J., J. Miller and R. Palmer (1993): "Behavior of Trading Automata in a Computerized Double Auction Market," in D. Friedman and J. Rust (Eds.), *The Double Auction Market: Institutions, Theories, and Evidence*, Addison Wesley. Chap. 6, pp.155–198.
27. Rust, J., J. Miller, and R. Palmer (1994), "Characterizing Effective Trading Strategies: Insights from a Computerized Double Auction Market," *Journal of Economic Dynamics and Control*, Vol. 18, pp. 61–96.
28. Shachat, J. and J. Swarthout (2002), "Procurement Auctions for Differentiated Goods," IBM Watson Research Center Working Paper.
29. Smith, V. L. (1991), *Papers in Experimental Economics*, Cambridge: Cambridge University Press.
30. Smith V. L., G. L. Suchanek, and A. W. (1988), "Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets," *Econometrica* **56(6)**, pp. 1119–1152.
31. Tay, N., and S. Linn (2001), "Fuzzy Inductive Reasoning, Expectation Formation and the Behavior of Security Prices," *Journal of Economic Dynamics and Control* **25**, pp. 321–361.
32. Varian, H. R. (1988), "Effect of the Internet on Financial Markets," School of Information Management and Systems, University of California, Berkeley.
33. Wellman, M., A. Greenwald, P. Stone, and P. Wurman (2002), "The 2001 Trading Agent Experiment," in *Proceedings of Fourteenth Innovative Applications of Artificial Intelligence Conference*, pp. 935–941.