

Finite automata games: basic concepts.

OLIVEIRA, F.S.

2014

© IGI Global. All rights reserved. This is the accepted manuscript version of the above chapter. The published version of record is available to purchase from the publisher's website:
<https://doi.org/10.4018/978-1-4666-5202-6.ch088>

Finite Automata Games – Basic Concepts

Editor(s) Name(s): John Wang

Book Title: **Encyclopedia of Business Analytics and Optimization**

Author(s) Name(s): Fernando S. Oliveira

Chapter Title: Finite Automata Games – Basic Concepts

Fernando S. Oliveira*

Operations Management and Decision Science Department, ESSEC Business School, 100 Victoria Street, National Library Building, #13-02, Singapore, 188064, Tel: +65 –68357865, Email:Oliveira@essec.edu

Summary:

In this chapter we review the basic concepts on automata games, including best response, inference, equilibrium and complex system dynamics. We describe how the concept of Nash equilibrium is used to analyze the properties of automata systems and discuss its limitations. We explain why we think the topics of automata inference, the modeling of evolving automata, and the analysis of the relationship between emotions and reason, are interesting areas for further research.

Classification: B: Business Process Optimization.

(*Corresponding author)

INTRODUCTION

Automata based systems have been used extensively in complex business modeling, for example, to represent Markov systems (e.g., Stewart et al., 1995; Uysal & Dayar, 1998; Gusak et al., 2003; Fuh & Yeh, 2001; Sbeity et al., 2008), in the development of classification systems (Gérard et al., 2005), in the analysis of commuters behavior (van Ackere & Larsen, 2004), to design electricity markets (Bunn & Oliveira, 2007, 2008), in the planning of real-options (Oliveira, 2010a), to study human-computer interaction (Gmytrasiewicz & Lisetti, 2002; Altuntas et al., 2007; Kim et al., 2010; Muller et al., 2013), to represent the relationship between emotions and reason (Oliveira, 2010c), in devising product differentiation strategies (Oliveira, 2010b), and in forecasting and production control (Liu et al., 2011).

The analysis of the behavior of such systems is very often based on the concepts of game theory, such as Nash equilibrium (e.g., Fudenberg & Tirole, 1991). The Nash equilibrium is a powerful tool for analyzing industries where there are strategic interdependences between players. However, it does not explain the process by which decision makers acquire equilibrium beliefs, failing to determine a unique equilibrium solution in many games, and, therefore, failing to predict, or prescribe, *rational behavior* (e.g., van Huyck et al., 1990; Samuelson, 1997; Fudenberg & Levine, 1998).

In games with multiple equilibria the Nash equilibrium fails to predict the players' behaviors. In this case, empirical studies (e.g., Roth & Erev, 1995) have shown that models of bounded rationality predict better than the Nash equilibrium does how people, organizations and markets behave (at least in the short run). A first attempt from the game theory literature to address this issue was to refine the concept of Nash equilibrium by including additional criteria. First, a player does not choose dominated strategies (Fudenberg

& Tirole, 1991: 8). Second, choices in information sets not in the equilibrium path must be optimal choices (in order to avoid non-credible threats). This is called the rationalizability criterion (Bernheim, 1984; Pearce, 1984). However, the problem with equilibria selection still exists as different refinements select different equilibria. Furthermore, rationalizable strategies may be too demanding as they assume common knowledge of rationality.

Therefore, in order to model complex games, possibly with multiple equilibria, computer models which incorporate boundedly rational players are used as a mechanism for inductive equilibrium selection, and to test the validity of the perfect-rationality predictions. This methodological jump from perfect-rationality to bounded rationality has theoretical and philosophical implications. It corresponds to a switch from a “normative theory” to a “positive theory.” The normative theory prescribes what each player in a game should do in order to promote his interests optimally (von Neumann & Morgenstern, 1953; van Damme, 1991: 1), whereas the positive theory describes how agents actually decide, as this line of research tries to understand how people and institutions behave (e.g., Samuelson, 1997: 3).

Simon (1972) was the first to emphasize the need to model bounded rationality in order to capture human and organizational behavior: see Sent (2004) for a review of Herbert Simon’s work. As Aumann (1997) explains, people and organizations use “rules of thumb” that they learned from experience when acting. In other words, people do not optimize even in simple decision problems, e.g., Salant (2011). This argument underlines the need to model the opponent’s behavior, which was formalized in Rubinstein (1986, 1998) using finite automata - see Hopcroft & Ullman (1979) or Cecherini-Silberstein et al. (2012) for an introduction to automata theory. In order for inference of the opponents’ strategic behavior to be possible some rules need to regulate the definition of strategies. Rubinstein (1986) proposed the finite automaton as a tool to model an agent’s behavior. Salant (2011) has used automata to implement choice rules. An automaton is a decision rule, or a strategy, consisting

of a finite set of states, a transition function (that defines the rules of transition between states) and a behavioral function (defining an agent's behavior in each state of the automaton). Rubenstein suggested that repeated games with finite automata could capture a player's bounded rationality (considering automata with a bounded number of states). At the same time, the introduction of finite automata constrains the type of strategies played: only regular strategies are admissible (i.e., given the same input, a player reacts always in the same manner). It is noteworthy, however, that long before Rubinstein had proposed the automata game, Schreider (1964) presented the formalism of dynamic programming to solve discrete deterministic problems using finite automata and introduced its possible application to game theory. In automata theory there are four major central issues: the complexity of computing the best response automaton, the equilibrium in automata games, automata inference and, finally, the dynamics problem.

THE FINITE AUTOMATA GAME

An automata game in the extensive form is a 5-tuple

$$G = \left(N, \{Z^i\}_{i=1}^N, \{u^i\}_{i=1}^N, \{Q^i\}_{i=1}^N, \{\Sigma^i\}_{i=1}^N \right). N \text{ denotes the number of players. } Z^i \text{ represents a}$$

finite non-empty set of possible outcomes of the game, and each $z^i \in Z^i$ is a function of the actions of each player, $z^i = z(a^i, a^{-i})$, where $a^i \in \Sigma^i$ represents an action of player i and $a^{-i} \in \Sigma^{-i}$ represents his opponents' actions. The outcomes of the game represent the information received by each player at the end of every stage. This information, or outcome, is a function of the actions of each player in the stage game, and it is different for each one of the players, as each one only knows the outcome of his own actions. $u^i = u(z^i)$ represents the utility function of player i , i.e., it is the payoff a player i perceives to have received from his

action, given the perceived outcome. Q^i stands for a finite non-empty set of internal states of player i . Σ^i is a non-empty set of all possible actions of player i . The automata game G is an extensive form game where each player evolves a certain decision rule that may change at a certain iteration of the game. This decision rule, the automaton A^i , defines how a player reacts to the outcomes received from the environment.

A finite automaton used by the player i is a 5-tuple $A^i = (Q^i, q_0^i, \Sigma^i, \delta^i, \lambda^i)$ in which: Q^i is a finite non-empty set of internal states; q_0^i is the initial internal state; Σ^i is the set of all the possible actions; δ^i is a transition function $(\delta^i : Q^i \times Z^i \rightarrow Q^i)$ and λ^i is a behavioral function $(\lambda^i : Q^i \rightarrow \Sigma^i)$ associating an action to each possible internal state. At stage one each player i plays $\lambda^i(q_0^i)$. At a stage $t \geq 1$, after each player executing his actions with an outcome $z_t^i = \lambda^i(a^i, a^{-i})$, each automaton A^i moves from the state q_t^i to the state $\delta^i(q_t^i, z_t^i)$. Then each player i chooses a new move $\lambda^i(q_{t+1}^i)$.

In an automata game, as Mor et al. (1996) put it, a player engages in three tasks at the same time: to define the strategy to play the game, to learn the opponents' strategies, and finally if the other players are also learning, to influence the opponents' beliefs. This line of research assumes that some of the players are not perfectly rational and that their limitations may be learned and exploited by the other players in the game.

COMPUTING THE BEST RESPONSE AUTOMATON

The complexity of computing the best response automaton aims to capture an agent's rationality and the cost of operating an automaton (e.g., Neyman, 1998). To measure this complexity has become a central issue in automata game theory. Rubinstein (1986) and

Neyman (1998) used the number of states in an automaton as a measure of the agent's complexity. Moreover, as the behavior and the transition functions are costly to operate, the decision-maker reduces these costs by minimizing the number of states in his automaton. Therefore, a player has profit maximization as his main objective and lexicographically he minimizes the number of states in the automaton (i.e., given two automata with the same expected profit the agent chooses the simplest one).

Banks and Sundaram (1990) generalized this measure of complexity in order to include the costs of monitoring the opponent's behavior. They argued that one needs to measure the "transitional complexity" in order to capture the full complexity of an automaton. Additionally, they compared the complexity of two automata taking into account both the number of states and the number of transitions between them. Banks and Sundaram, as well as Piccione (1992), have modeled best response in automata games within a discounted infinite stage model. In this context, Piccione proved the equivalence between the policy derived by best response and the stationary policy computed by dynamic programming.

Gilboa (1988) defined the problem faced by a player using a finite automaton: to choose the best-response automaton, given the choices of the other players. He defined Nash equilibrium of the automata game as the choice of an automaton, for each player, such that no player can increase his payoff by unilaterally changing his automaton. He focused on the complexity of *computing the best response automaton* and not on the complexity of implementing it. This is a very strong criticism of the previous literature. He argued that limitations on the number of states do *not* capture bounded rationality, which is a restriction on the capabilities of an agent to *design* a strategy, while the constraint on the number of states only captures a limitation on the *implementation* of a strategy. Furthermore, also on this topic, Papadimitriou (1992) has shown that there is a trade-off between the designing and the

implementation complexities. In other words, the problem of computing the best response automaton with a constraint on the possible number of states is *NP-complete* whereas, if no constraint is given, this problem can be solved in polynomial time.

Gilboa's work gave rise to a series of papers on the complexity of computing the best response automaton. Ben-Porath (1990) has analyzed the calculation of the best response automaton under uncertainty (assuming that the opponent has a set of possible automata to use during the game). He proved that, when there is uncertainty concerning the automata the other players use (and he studied a game against nature), the problem of computing the best response automaton is *NP-Complete*, meaning that there is no polynomial algorithm to solve it. Finally, a probabilistic automaton that chooses actions randomly can also capture uncertainty. Freund et al. (1995) demonstrated that, in this case, the complexity of finding the best response automaton against a random opponent is equivalent to the complexity of the deterministic case. Gossner and Hernandez (2003) have also studied the complexity of coordinated play in automata games, allowing for a few mismatches, proving that for any sequence of actions there is an automaton which achieves coordination with a ratio close to 1.

On this same topic, Lehrer and Solan (2006) have addressed the concept of excludability in the context of automata games with vector payoffs, analyzing the conditions under which, in a two-players game, one of the players has a strategy that ensures that if the opponent uses an automaton the average payoff in the long-run is bounded away from the set.

A related issue regards the conditions under which such complex systems can emerge from evolution rather than from optimal design. Kilani (2007) has proposed an evolutionary model of finite automata showing that, in the case of long interactions, cooperation emerges as an equilibrium strategy. Oliveira (2010c) has analyzed how automata can be used to model emotions in agent-based systems in which agents co-evolve endogenous rules of behavior. He proved that best response is not sufficient to define complete and consistent rules of behavior

and that, instead, emotions are necessary to enable the agent to improve his behavior when adapting to others.

EQUILIBRIUM IN AUTOMATA GAMES

The computation of the equilibrium of the automata game is the second major issue addressed in the literature. Abreu and Rubinstein (1988) have shown that automata games have Nash equilibria. This proof is available in a setting where there are two players that play a one-shot game with complete information and, therefore, assuming that the players *always use the same automaton*. Furthermore, Abreu and Rubinstein proved that in a Nash equilibrium of an automata game (with two players), each player uses a finite automaton with an equal number of states. Since the automata are finite, the game eventually reaches a *cycle* where it repeats the pairs of states played. This introduces a partition of state pairs into those belonging to a cycle and those only played at the beginning of the game. The states of a player's automaton that appear in a cycle are all distinct: the other states appear in the beginning of the play and are never repeated. Therefore, in equilibrium, there is a one-to-one correspondence between the stage-game actions of the two automata.

However, even in such a simple setting the use of the Nash equilibrium as a tool of analysis is an issue. Gilboa (1988) criticizes the concept of Nash equilibrium as the number of players in a game is usually not known. In addition, Gilboa and Zemel (1989) do not support the use of the Nash equilibrium concept in the analysis of automata games as the problem of determining its existence or its uniqueness is *NP-Hard*. Finally, Babichenko (2010) has analyzed the game in which every player knows only his own payoff function (with uncoupled automata) showing that, in this case, there are pure Nash equilibria. It should be noted that in Oliveira (2010 b, c) the players infer or compute the best response against an

aggregated automaton of the automata used by all his opponents and that the player only observes his own payoff function and, for this reason, these are also uncoupled automata, as defined in Babichenko (2010).

AUTOMATA INFERENCE

In an automata game, a player faces not only the control problem (to optimize his automaton against a given opponent) but also an inference problem (to learn the rule or rules of behavior used by his opponent or opponents).

There are two major branches in the automata learning literature that address the inference problem (e.g., Angluin, 1987; Gold, 1978; Oliveira, 2010b): active and passive learning. A learning algorithm tries to infer the automaton generating a stream of data. Active learning is the inference problem faced by a player who has the ability to influence the input generation process, i.e., he is able to select the inputs for the automaton generating the data. In a passive learning problem, a player has no control over the inputs supplied to an automaton.

Thus, the actions of the learning player, in an active learning algorithm, affect the output of the automaton generating the data. Hence, each player faces an active learning problem. In this case, Angluin (1987) has proved that a learning algorithm that attempts to infer the behavior of an automaton, provided with counterexamples by a benevolent teacher, and with the possibility of controlling the inputs to the target automaton, can learn the automaton's correct structure in polynomial time. Furthermore, Angluin also defines the *minimum requirements* to learn the best possible rational model of an automaton's behavior: *completeness* and *consistency*. The completeness requirement implies that a player holds a forecast for every action in every state of the automaton model. The consistency requirement

implies that a player holds a correct model of the automaton he is inferring. A model is correct if, in a certain state, it does not forecast different transitions for the same action. Thus, the completeness and consistency requirements enable a player to infer a model where the transition and behavioral functions are complete and have no contradictions.

However, Angluin's algorithm is not satisfactory in modeling competitive games as in most game-theoretical models there is a very high cost in experimenting with opponents. This was first realized by Carmel and Markovitch (1996, 1999). They developed an algorithm to infer the automaton used by an opponent. The player inferring his opponent's behavior uses as only source of information his own experiences in interacting with that given opponent. Oliveira (2010b) has modeled quasi-perfect learning which studies the behavior of individuals or organizations that can get the correct answer to any of their queries (as they are very knowledgeable) but are not provided with counterexamples (as no benevolent teacher exists to guide learning in such environments) proving that even in such conditions the correct automaton can be learned if the algorithm uses a long enough planning horizon.

AUTOMATA DYNAMICS

As recognised by Rubinstein (1998) the automata game, as formulated in this literature, is a one-shot game in which the players cannot alter the automaton and where dynamic aspects are not considered.

However, recently, there have been several attempts (coming from the artificial intelligence area) to incorporate dynamic issues into automata games. Carmel and Markovitch (1999) developed an algorithm that enables a player to infer a model of his opponent by interacting with him. This algorithm incorporates an endogenous exploration

mechanism that enables a player to plan his actions in advance (an agent modifies his behavior and learns from the interaction with his opponent). Thus, a player attempts to profit from his knowledge and, at the same time, tries to improve the inferred model when his predictions are incorrect. Nonetheless, this research still assumes a stationary opponent, i.e., it assumes that the opponent does not change his automaton.

In the study of the relationship between emotions and reason Oliveira (2010c) has proposed a setting in which the agents are allowed to change their automaton by using several possible operations, including emotions and best-response, allowing all the agents to readjust to each others' behaviors during the adaptation process. This model considers that the players are evolving over time in a non-stationary environment.

CONCLUSIONS AND FUTURE TRENDS

The modeling of automata enables the design of an agent that infers his opponent's behaviors and exploits this knowledge by creating better strategies to play the game. This methodology seems a promising avenue in creating models of systems of companies that use rules of behavior, enabling the analysis of behaviors that are not planned but, instead, evolve over time.

Automata based models have been used to capture human-machine cooperation systems, allowing the development of human machine cooperation in dynamic settings (Gmytrasiewicz & Lisetti, 2002; Altuntas et al., 2007; Kim et al., 2010). This topic seems to be of interest for future research as with the increasing automation of the working place, and of society in general, the study of the human-machine interaction will, certainly, be central to the development of the society of the future.

The use of automata theory to represent the relationship between emotions and reason in the construction of rules of behavior seems to be an interesting research direction (Oliveira, 2010c). We envision that experimental studies with people to test these theories could lead to important discoveries on the interaction between emotion and reason in game playing.

Automata inference, both in the passive and active forms, seems to be a promising topic as well, as it has the potential to transform data on the behavior of people and organizations into decision rules that explain their internal decision processes. This knowledge is of interest to model competitor's behavior and it is very relevant for the management of complex organizations, possibly allowing a better structuring of decision systems.

REFERENCES

Abreu, D., & Rubinstein, A. (1988). The structure of Nash equilibrium in repeated games with finite automata. *Econometrica*, 56(6), 1259-1281.

Altuntas, B., Wysk, R. A., & Rothrock, L. (2007). Formal approach to include a human material handler in a computer-integrated manufacturing (CIM) system. *International Journal of Production Research*, 45(9), 1953-1971.

Angluin, D. (1987). Learning regular sets from queries and counterexamples. *Information and Computation*, 75, 87-106.

Aumann, R. J. (1997). Rationality and bounded rationality. *Games and Economic Behavior*, 21, 2–14.

Babichenko, Y. (2010). Uncoupled automata and pure Nash equilibria. *International Journal of Game Theory*, 39, 483-502.

Banks, J. S., & Sundaram, R. K. (1990). Repeated games, finite automata, and complexity. *Games and Economic Behavior*, 2, 97-117.

Ben-Porath, E. (1990). The complexity of computing a best response automaton in repeated games with mixed strategies. *Games and Economic Behavior*, 2, 1-12.

Bernheim, D. (1984). Rationalizable strategic behavior. *Econometrica*, 52, 1007-1028.

Bunn, D. W., & Oliveira, F. S. (2007). Agent-based analysis of technological diversification and specialisation in electricity markets, *European Journal of Operational Research*, 181(3), 1265-1278.

Bunn, D. W., & Oliveira, F. S. (2008). Modeling the impact of market interventions on the strategic evolution of electricity markets, *Operations Research*, 56(5), 1116-1130.

Cecherini-Silberstein, T., Coornaert, M., Fiorenzi, F., & Schupp, P. E. (2012). Groups, graphs, languages, automata, games and second-order monadic logic. *European Journal of Combinatorics*, 33, 1330-1368.

Carmel, D., & Markovitch, S. (1996). Learning models of intelligent agents. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*. Portland, Oregon, 62-67.

Carmel, D., & Markovitch, S. (1999). Exploration strategies for model-based learning in multi-agent systems. *Autonomous Agents and Multi-agent Systems*, 2, 141 – 172.

Freund, Y., Kearns, M., Mansour, Y., Ron, D., Rubinfeld, R., & Schapire, R. E. (1995). Efficient algorithms for learning to play repeated games against computationally bounded adversaries. In *proceedings of the 36th IEEE Symposium on the Foundations of Computer Science*, 332-341.

Fudenberg, D., & Tirole, J. (1991). Game theory. The MIT Press, Cambridge, Massachusetts.

Fudenberg, D., & Levine, D. K. (1998). The theory of learning in games. The MIT Press, Cambridge, Massachusetts.

Fuh, C.-F., & Yeh, Y.-N. (2001). Random perturbation in games of chance. *Studies in Applied Mathematics*, 107, 207-215.

Gérard, P., Meyer, J.-A., & Sigaud, O. (2005). Combining latent learning with dynamic programming in the modular anticipatory classifier system. *European Journal of Operational Research*, 160, 614-637.

Gilboa, I. (1988). The complexity of computing best-response automata in repeated games. *Journal of Economic Theory*, 45, 342-352.

Gilboa, I., & Zemel, E. (1989). Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior*, 1, 80 – 93.

Gmytrasiewicz, P. J., & Lisetti, C. L. (2002). Emotions and personality in agent design and modeling, in *Game Theory and Decision Theory in Agent-Based Systems*, Ed. S. Parsons, P. J. Gmytrasiewicz, and M. Wooldridge, 81-95.

Gold, E. M. (1978). Complexity of automaton identification from given data. *Information and Control*, 37, 302-320.

Gossner, O., & Hernandez, P. (2003). On the complexity of coordination. *Mathematics of Operations Research*, 28(1), 127-140.

Gusak, O., Dayar, T., & Fourneau, J.-M. (2003). Lumpable continuous-time stochastic automata networks. *European Journal of Operational Research*, 148, 436-451.

Hopcroft, J. E., & Ullman, J. D. (1979). Introduction to automata theory, languages and computation. Addison-Wesley, Mass.

Kilani, M. (2007). Evolution and the complexity of finite automata. *International Game Theory Review*, 9(4), 731-743.

Kim, N., Shin, D., Wysk, R. A., & Rothrock, L. (2010). Using finite state automata (FSA) for formal modeling of affordances in human-machine cooperative manufacturing systems.

International Journal of Production Research, 48(5), 1303-1320.

Lehrer, H., & Solan, E. (2006). Excludability and bounded computational capacity.

Mathematics of Operations Research, 31(3), 637-648.

Liu, H., Zabinsky, Z. B., & Kohn, W. (2011). Rule-based forecasting and production control system design utilizing a feedback control architecture. *IIE Transactions*, 43, 143-152.

Mor, Y., Goldman, C. V., & Rosenschein, J. S. (1996). Learn your opponent's strategy (in polynomial time)!. G. Weiss and S. Sen (eds.). *Lecture Notes in Artificial Intelligence*, 1042, Springer Verlag.

Muller, B., Bohn, F., Drebler, G., Groeneveld, J., Klassert, C., Martin, R., Schluter, M., Weise, H., & Schwarz, N. (2013). Describing human decisions in agent-based models – ODD +D, an extension of the odd protocol. *Environmental Modeling & Software*, 48, 37-48.

Neyman, A. (1998). Finitely repeated games with finite automata. *Mathematics of Operations Research*, 23(3), 513-552.

Oliveira, F. S. (2010a). Bottom-up design of strategic options as finite automata.

Computational Management Science, 7(4), 355-375.

Oliveira, F. S. (2010b). Limitations of learning in automata-based systems. *European Journal of Operational Research*, 203(3), 684-691.

Oliveira, F. S. (2010c). Modeling emotions and reason in agent-based systems. *Computational Economics*, 35, 155-164.

Papadimitriou, C. H. (1992). On players with a bounded number of states. *Games and Economic Behavior*, 4, 122 – 131.

Pearce, D. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52, 1029-1050.

Piccione, M. (1992). Finite automata equilibria with discounting. *Journal of Economic Theory*, 56, 180-193.

Roth, A. E., & Erev, I. (1995). Learning in extensive-form games: experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior*, 8, 164-212.

Rubinstein, A. (1986). Finite automata play the repeated prisoner's dilemma. *Journal of Economic Theory*, 39, 83-96.

Rubinstein, A. (1998). Modeling Bounded Rationality. MIT Press.

Salant, Y. (2011). Procedural analysis of choice rules with applications to bounded rationality. *American Economic Review*, 101(2), 724-748.

Samuelson, L. (1997). Evolutionary games and equilibrium selection. The MIT Press.

Sbeity, I., Brenner, L., Plateau, B., & Stewart, W. J. (2008). Phase-type distributions in stochastic automata networks. *European Journal of Operational Research*, 186, 1008-1028.

Schreider, Y. A. (1964). Automata and the problem of dynamic programming. *Problems of Cybernetics*, 5, 33-58.

Sent, E.-M. (2004). The legacy of Herbert Simon in game theory. *Journal of Economic Behavior & Organization*, 53, 303-317.

Simon, H. A. (1972). Theories of bounded rationality, in Macguire and Radner (Eds.), *Decision and Organisation*, North-Holland.

Stewart, W. J., Atif, K., & Plateau, B. (1995). The numerical solution of stochastic automata networks. *European Journal of Operational Research*, 86, 503-525.

Uysal, E., & Dayar, T. (1998). Iterative methods based on splitting for stochastic automata networks. *European Journal of Operational Research*, 110, 166-186.

van Ackere, A., & Larsen, E. R. (2004). Self-organizing behavior in the presence of negative externalities: A conceptual model of commuter choice. *European Journal of Operational Research*, 157, 501-513.

van Damme, E. (1991). Stability and perfection of Nash equilibria. Second edition, Springer-Verlag.

van Huyck, J. B., Battalio, R. C., & Beil, R. O. (1990). Tacit coordination games, strategic uncertainty, and coordination failure. *The American Economic Review*, 80(1), 234-248.

von Neumann, J., & Morgenstern, O. (1953). Theory of games and economic behavior. Third Edition. Princeton: Princeton University Press.

KEY TERMS & DEFINITIONS

Active Learning: during the process of automata inference the player needs to estimate the internal states of the process observed (together with the respective behavioral and transition functions) by observing the sequence of inputs and outputs. If the player *has control* over the inputs provided to the automaton generating the data we have an active learning process.

Automaton: it is a decision rule, or a strategy, consisting of a finite set of states, a transition function (that defines the rules of transition between states) and a behavioral function (defining an agent's behavior in each state of the automaton).

Automata Game: it is a game in which the players' strategies are rules of behavior encapsulated in an automaton that describes how the agent behaves in each state and how he reacts to changes in his environment. The automata employed by the agents may have been

the result of an optimization procedure at the start of the game or have evolved through learning over time.

Automata Inference: the player attempts to learn the automata that can better describe the rules of behavior used by the process he interacts with or employed by a given opponent or set of opponents. Automata learning includes the observation of the inputs and outputs produced by a given system and, from these data, the estimation of the states in the automaton and respective behavioral and transition functions. There are two types of inference processes, active and passive learning, depending on how much control the player has on the inputs to the system (or opponent) he is observing.

Best Response Automaton: it is the rule of behavior (composed by the internal states and by the behavior and transition functions) that maximizes the expected utility received by the automaton when employed against a given opponent's automaton.

Nash Equilibrium of the Automata Game: it is a state of the automata game in which the choice of automaton by each player is such that no player can increase his payoff by unilaterally changing his automaton.

Passive Learning: it is the process used by an agent to learn the automaton representing the behavior of a different system when he has *no control* over the inputs supplied to an automaton. In this case the agent is a passive observer of the behavior of the system without interacting with it.