

CONNECTING ADAPTIVE BEHAVIOUR  
AND EXPECTATIONS IN MODELS OF INNOVATION:  
THE POTENTIAL ROLE OF ARTIFICIAL NEURAL  
NETWORKS

MURAT YILDIZOĞLU<sup>1</sup>

**Abstract.** In this methodological work I explore the possibility of explicitly modelling expectations conditioning the R&D decisions of firms. In order to isolate this problem from the controversies of cognitive science, I propose a *black box* strategy through the concept of “internal model”. The last part of the article uses artificial neural networks to model the expectations of firms in a model of industry dynamics based on Nelson & Winter (1982).

**Classification Codes.** D83, 03.

*“And one of the deepest, one of the most general functions of living organisms is to look ahead, to produce future as Paul Valéry puts it.” (François Jacob, 1982, p. 66).*

## 1. INTRODUCTION

The purpose of R&D investment is to carry out an innovation which is a potential source of competitiveness for firms. When a firm invests in R&D, it becomes involved in a dynamic process which rests on an *a priori* belief in technological progress. The primary motivation of firms to invest in R&D is linked to this expectation on the existence of a technological change process that firms try to exploit in order to increase (or to maintain) their competitiveness. Thus R&D investment

---

*Keywords and phrases:* Neural networks, genetic algorithms, bounded rationality, learning, expectations, innovation dynamics.

<sup>1</sup> IFREDE-E3I, Université Montesquieu Bordeaux IV, Avenue Leon Duguit, F-33608 Pessac, France. e-mail: [yildi@montesquieu.u-bordeaux.fr](mailto:yildi@montesquieu.u-bordeaux.fr)  
<http://yildizoglu.montesquieu.u-bordeaux.fr/>

© EDP Sciences 2002

corresponds to a decision which is turned to the future and, as a consequence, which involves an expectation dimension.

At the same time, R&D decisions are characterized by a strong uncertainty concerning the return on investment. This uncertainty is stronger for R&D investment than for other types of investment. Indeed innovations often result from what Simon (1958) calls “nonprogrammed decisions”, that is situations where the alternatives of choices must be discovered by firms and the connections between choices and consequences are imperfectly known. It is the reason why R&D decisions are generally associated to uncertainty in the sense of Knight (1921). This uncertainty strongly limits the ability of firms to form expectations about the return on their R&D investment. In this context, firms must be able to improve through experience their perception of the relationships between R&D investment and competitiveness and to adapt accordingly their R&D decisions. Thus firms determine their R&D investment through the combination of adaptation with expectation of potential return on R&D. Both dimensions generally coexist except in the extreme situations where there is no uncertainty or where the uncertainty is totally radical so that firms can not form any expectation.

In order to emphasize the role of uncertainty in R&D decisions, we should distinguish between two types of uncertainty:

- the technological uncertainty concerns the connection between R&D and innovation. It depends on the nature of innovation (radical or incremental) and on the potentialities of the technology which is exploited by firms. This uncertainty influences the occurrence and the time of innovation as well as the technological performances associated to the innovation;
- the market uncertainty rather affects the link between R&D investment and competitiveness of firms (in terms of profit or market share). The impact of innovation on the competitiveness of firms does not only depend on technological factors but also on the evolution of demand and on the behaviors of competing firms. This uncertainty is mainly linked to the environment in which R&D is carried out, especially when the industry persistently is out of equilibrium.

On the basis of this distinction between the sources of uncertainty, I can argue that the nature and the degree of uncertainty in R&D decisions depend on the technology and on the market structure. In this perspective, uncertainty is not systematically associated to radical innovation but to innovation in general. For instance in an oligopolistic structure, a firm that carries out R&D activities in order to realize incremental innovations has to cope with the market uncertainty which comes from the behavior of its rivals. In an equilibrium approach this problem is overcome since every firm is supposed to perfectly anticipate the equilibrium decisions. Otherwise this uncertainty strongly influences the impact of innovation upon the competitiveness of the firm. Another case could be a monopolistic situation with radical innovation. In that case, uncertainty would come mainly from technological factors (unless unpredicted variations of demand occur). These examples clearly show that the combination of expectation with adaptation

in the presence of uncertainty is a general problem and it is relevant for all types of innovation (incremental or radical).

I propose to explore a way of modelling of the determination of the level of R&D investment of firms. As a consequence, this article will not tackle the decision of being an innovator or not, nor the adoption of a new technology. I will exclude these decisions and focus on the situations where firms invest in internal R&D in order to produce an innovation. In that case the problem is to determine the level of R&D investment. My interest is to analyze how expectation and adaptation can be combined in the modelling of R&D investment rules. These dimensions are generally split up in the literature: rational expectations are assumed in neoclassical models, whereas alternative approaches (institutional and/or evolutionary) generally adopt a purely adaptive representation.

The concept of equilibrium with rational expectations or perfect foresight, which is generally used in standard models, leads to a systematic coordination among agents and among periods. In that sense, the role of rational expectations is mainly to insure intertemporal coordination. The problem with this approach is that it does not deal with how agents form and adapt their expectations. Under the equilibrium assumption there is no need for adaptation since the expectations correspond to the objective distribution of outcomes. In this context, agents do not have to exploit observations and past experiences in order to improve their predictions and to adapt the model of the economy they use to form expectations. However this issue is particularly relevant for R&D decisions since they are guided by the agents' vision of the technological change process. Also, the neoclassical models generally deal with risk (*v.s.* uncertainty) and they do not cope with the formation and the adaptation of expectations. As such, I exclude them from the scope of this article.

In the next section, I focus on models based on adaptive decision rules. I show that recent models of innovation developed in the evolutionary framework tend to overlook the role of expectations in innovation decisions. In this class of models, the decision rules are purely adaptive, and as a consequence, they do not take into account the impact upon investment of the expectations of firms concerning the relationships between R&D and competitiveness.

Section 3 is devoted to the modelling trade-off between expectation and adaptation. I argue that in the presence of uncertainty, expectations reflect the existence of an internal model of the economy that firms use to make simulations about the possible outcomes of their decisions. This internal model is specific to firms and is adapted over time according to observations and experiences. The arguments of this section show that artificial neural networks (ANN) appear as a natural candidate for modelling the presence of such an internal model.

This assumption is tested in Section 4 where an evolutionary model of industry dynamics is developed. This model builds upon Yildizoglu (2001) (see the presentation of this model at p. 209) where some firms (the GenFirms) explore the R&D strategy (investment rate) through a genetic algorithm. In this first, deliberately crude, version of the model I assumed that the only possibility for the firms to learn the fitness of each rule was to effectively use it during a learning period. In

the model of the fourth section, the fitness of each rule is evaluated on the base of firm's expectations provided by a firm-specific neural network. The results from different assumptions about learning are presented: Nelson and Winter kind fixed-rule based behaviour (NWFirms), Genetic Algorithm (**GA**) based learning where each rule is effectively used (GenFirms) and GA based learning with expectations based on a neural network (NNGenFirms).

## 2. UNCERTAINTY AND ADAPTIVE BEHAVIOR

When one takes into account the complexity of the innovation process and the inherent uncertainty (technological and/or market), the assumption of agents possessing the true model of the economy becomes very costly in terms of realism. This assumption is of course critical for models based on rational expectations.

Evolutionary modelling of R&D decisions has rejected this assumption from the beginning (see Nelson and Winter, 1982 and Hodgson, 1994 for a historical account). As an alternative, the bounded rationality concept of Simon and adaptive behavior based on decision rules have been placed in the center of evolutionary models (see Simon, 1982).

Alternatives to the unifying rational expectations framework lead to a great diversity because the modelling of bounded rationality entails the modelling of agents' learning process and an empirically and conceptually founded unanimity about this learning process does not exist for the time being (see Dennett, 1998). Simon's propositions have consequently been introduced in economic models in a progressive and diversified way. Concerning the modelling of R&D decisions, the persistent form of this formalization has been centered around decisions rules and their modification as a consequence of learning.

### 2.1. FROM BOUNDED RATIONALITY TO DECISION RULES

The original model of Nelson and Winter (1982, ch.15) introduces a very simple kind of R&D investment rule: each firm is supposed to invest a constant fraction of its capital stock in R&D. Even if the merit of this model is to propose a first application of the concept of bounded rationality to innovation decisions, it is limited by a very simplistic version of this concept. In the end, as Silverberg and Verspagen (1999) assert: "While there is technological learning at the economy-wide level, firms themselves are completely unintelligent, since they operate according to given search and investment rules that cannot be modified as a result of experience." This is of course very far from Simon's initial propositions.

Winter (1984) provides a first attempt in introducing adaptive R&D decision rules. In this model, the R&D-capital ratios are adapted in order to reach a satisfactory profit rate which is given by the capital weighted industry average profit rate. If the current profit of the firm is less than satisfactory, the firm progressively adjusts its R&D-capital ratio towards the industry average ratio. This adaptive rule corresponds more to blind imitation (of the average behavior) than to the positive result of a learning process. Moreover this rule assumes

<i>Population at date t</i>	<i>Fitness: f(x)=x<sup>2</sup></i>	<i>Expected number: f(x)/Mean<sub>t</sub></i>	<i>Effective number proportionally drawn</i>	<i>New population</i>
1) 00011 (=3) 2) 01100 (=12) 3) 10111 (=23)	1) 9 2) 144 3) 529 Mean <sub>t</sub> =227	1) 0.0 2) 0.6 3) 2.4 Sum=3	1) 0 2) 1 3) 2 Sum=3	1) 10111 (=23) 2) 01100 (=12) 3) 10111 (=23)
<i>Crossover (3)-(2) at bit 3</i>				
1) 10111 (=23) 2) 01111 (=15) 3) 10100 (=20)	<i>Mutation (1)</i>			<i>Population at date t+1</i>
	1) 11111 (=31*) 2) 01111 (=15) 3) 10100 (=20)			1) 11111 (=31) 2) 01111 (=15) 3) 10100 (=20)
<p>Simple application of GAs to optimisation of the function <math>f(x)=x^2</math> over the interval 0-31. Integers are coded with five bits binary code: 00001=1, 11111=31. The example uses an initially random population of 3 members and the GA constructs a new population through selective reproduction, combination (crossover) and random experiments (mutation). In this schematic example, the GA attains the optimum (31) in one period. For each string, the crossover, its position and the partner, as well as mutation position are chosen randomly. The mutation bit simply switches its value: 0-&gt;1 or 1-&gt;0. <b>This process is controlled by: population size, bit-string size, probability of crossover and probability of mutation.</b></p>				

FIGURE 1. A simple example of genetic algorithm.

that each firm knows the distributions of capital stocks, profit rates and R&D ratios in order to be able to compute the weighted averages. Hence we have an informatively demanding but finally poor adaptive behavior behind this rule. Silverberg and Verspagen (1996, 1995) enrich this rule by letting firms choose their R&D investment as a fraction of profit or sales. Moreover the relative weight of each source of R&D financing (sales or profit) is adapted through imitation and random experiments (mutations) according to a satisficing rule. This emphasizes a second important dimension of the evolutionary modelling of learning process: random experiments or mutations. As a matter of fact imitation and mutation are two main dimensions of adaptive rules that will be introduced in subsequent models of innovation.

## 2.2. TOWARDS RICHER ADAPTATION: THE EMERGENCE OF A UNIFIED MODELLING PRINCIPLE?

Recent and more sophisticated models of technology dynamics progressively adopt mechanisms of rule adaptation more resolutely inspired by evolutionary algorithms. These algorithms are not introduced to represent faithfully the exact learning mechanism of agents but to just take into account, in the least ad hoc way, the presence of learning (see Marengo, 1992, for a precursory application of evolutionary algorithms to models of learning).

The models developed by Kwasnicki (see Kwasnicki and Kwasnicka, 1992; Kwasnicki, 1998) use a representation of the learning of firms already very close to genetic algorithms (see Fig. 1 and Goldberg 1991). Even if the R&D strategies of firms do not directly result from learning and adaptive processes, the result of this R&D (*i.e.* innovation) is modeled as the discovery and the effective use of new

routines: “The creative process is evolutionary by nature and as such its description should be based on a proper understanding of the hereditary information.” (Kwasnicki, 1998, p. 140) The learning of the firm does not directly concern its strategy of R&D investment (R&D is a constant function of capital stock) but the routines which enable it to produce better products. These routines are adapted through an evolutionary process based on recombination and mutation. The speed and the scope of this adaptation depend on the decisions of the firm (mainly on R&D investment). Given the emphasis on the modelling of innovation as the learning of new routines, these models open yet a little more the “black box” of the firm and, hence, their adaptive dimension is closer to the behavioral theory of the firm than to the innovation theory.

An even richer modelling of adaptive R&D strategies can be found in the work of Ballot and Taymaz. Even if R&D decisions are on a secondary plan in their first articles (see for example Ballot and Taymaz, 1997), one of their most recent articles is dedicated to R&D rules. In their initial articles they use controlled Genetic Algorithms to model the learning process of the firm as a result of incremental innovations. The efficiency of this learning is mainly determined by the spending of firm in human capital training. Even if these models are remarkable because they plainly adopt evolutionary algorithms, they are closer to Kwasnicki (1998) since the adaptation concerns the technology of firms and not their innovative strategies. Ballot and Taymaz (1999) is directly dedicated to the comparative modelling of R&D strategies. This model confronts an evolutionary adaptation mechanism, which is a Classifier System (CS), with other more ad hoc satisficing criteria: Informed behavior *à la* Nelson and Winter (1982), optimizing behavior based on statistical inference about the production function and Follower behavior which is pure imitation of the top 50% of firms in the industry.

A CS uses rules like “IF (Condition) THEN (Action)”. Each time an Action is chosen only if its Condition is fulfilled and if the rule has a strong value. This value is fixed by a credit assignment system and new rules are generated by a GA. Because of the condition part, the CS has the possibility to develop in parallel better sets of rules for different situations. Moreover, recent formulations of CS include a more direct reference to learning and expectations, through the inclusion of the prediction errors in the evaluation of the fitness (see XCS developed by Wilson, 1995).

In Ballot and Taymaz (1999), this CS continuously adapts the R&D-Sales ratio according to the information about the market share and the relative performance of the firm on the market. Then the GA generates new rules with a speed which is increasing in the general human capital of the firm. The comparison between these four rules clearly shows that the combination of evolutionary algorithms with empirical knowledge on the elements of R&D decisions can significantly limit the ad hoc nature of rule-based adaptive models.

The main idea of GA (and also of the CS) is the combination of good solutions in order to obtain better and even the best solutions. This combination is augmented by casual random experimenting. This simple idea is behind the Schemata Theorem of Holland (Goldberg, 1991; Mitchell, 1996).

Yildizoglu (2001) tests this idea in a very simple model based on Nelson and Winter (1982) model. This article makes rule-based firms compete with learning firms. The R&D and capital investment of the latter is modelled using a very simple genetic algorithm for each of them. Two main class of results show that we should not neglect learning if we desire to better understand the connection between innovative activity and industry dynamics. In the first place, results at the industry level show that industries with learning firms exhibit higher technological and allocative efficiency (welfare). In the second place, learning gives a competitive edge to firms benefiting from it: learning firms very quickly dominate the industry. Even this simple model shows that we cannot ignore learning since its presence has a tremendous impact on firms' and industry performance.

This simple mechanism of learning is also the basis of an emerging paradigm in Artificial Life studies. If one aims to explore the behavior of a whole system (an ecology) of adaptive agents, this can be done in the unifying frame of Complex Adaptive Systems (CAS) (Holland, 1996; 1998). This paradigm is actually far from providing a rich unified framework for models of economic decisions. In the meantime, evolutionary algorithms combined with a good empirical knowledge on economic decisions can offer a coherent strategy for adaptive models. This strategy could overcome the criticism concerning the excessive diversity of bounded rationality models by providing such a unified framework (see Sargent, 1995, for a summary, and Rubinstein, 1998 for alternatives that stick to an axiomatic approach and to the equilibrium assumption). Such a common framework would unfortunately not solve the main problem of adaptive models: because they refuse to assume that agents possess the true model of the economy in their head, these models go to the other extreme by considering agents without any model in their mind. In that case agents just "grope in the dark". We must hence find a convincing method for modelling expectations.

### 3. UNCERTAINTY AND MODELLING OF EXPECTATIONS

If one does not adopt a biologically over-determined vision of human behavior ("everything is coded in our genes"), the intentions of economic agents must be taken into account in the modelling of their decisions. Even with minimal knowledge, under uncertainty, agents look forward: a firm engages in R&D investments only if it assumes that there is potential technical progress and, moreover, that research activities can give a competitive edge over other firms through new production processes or products. R&D activity consequently follows from the anticipation of a positive impact of this activity on the prospects of the firm.

Unfortunately, it is very difficult to include these anticipations in models of R&D because we do not have a precise knowledge on how these expectations are actually formed by economic agents. The adaptive learning mechanisms of the preceding section include an expectational dimension but only for conditions already observed and rules already used by agents. Expectations are hence *hard-coded* into rules and learning about the environment cannot be distinguished from

the search for better rules (Dosi *et al.*, 1999). Consequently, these mechanisms do not explain how agents represent their environment in order to evaluate decisions even if they are not yet been tested. To overcome this problem, maybe we should accept to be unable to model the formation process of expectations and focus more on the effects of the simple presence of such a process. This black-box approach would leave aside a detailed representation of this process and thus overcome one of the major problems of Artificial Intelligence (Hofstadter, 1979; Dennett, 1998).

### 3.1. EXPECTATIONS AND AGENT'S INTERNAL MODEL

Instead of assuming that the agents know the exact model of the economy (rational expectations hypothesis), one can imagine that their decisions are guided by a more or less approximate model. This model summarizes the state of the agent's knowledge and evolves as a consequence of evolution of this knowledge. This approximate model can be called the internal model of the agent. It guides the decision process since it enables the agent to test the connections between the alternatives of choice and their consequences. The presence of such an internal model can reflect the intentionality of decisions.

Obviously, in this context, the concept of "model" must be understood in a very loose sense. More than a mathematical construction, it consists in a representation of the agent's perception of the environment: "In (...) situations [that are not sufficiently simple as to be transparent to human mind], we must expect that the mind will use such imperfect information as it has, will simplify and represent the situation as it can, and make such calculations as are within its powers" (Simon, 1976, p. 144). These calculations are "As if" experiments that enable the agent to evaluate the possible consequences of its decisions. In other words, before making a decision, the agent simulates the potential outcomes of different decisions by using its internal model. The output of these simulations yields the expectations of the agent.

Concerning the R&D decisions, the relevance of such a model could seem more problematic. These decisions are made by a meta-agent: the firm. A faithful description of the internal model of such a meta-agent is of course impossible, and in a certain sense meaningless. Fortunately, one does not need such a description in order to embody expectations in models of R&D. The effects of the intentionality can be studied simply by representing the presence of such a model instead of its exact structure. This is not necessarily a restrictive assumption since "the actions taken within organization need to be consistent; the frameworks within which they are embedded do not. (...) All is required is that the frameworks should fit where they touch" (Loasby, 1986, p. 51). In this case, one can easily assume that the firm bases its decisions on the connection that its internal model establishes between R&D and the relevant dimensions of the environment.

The agent compares the expectations resulting from the simulations with effective observations. If this experience questions the internal model, the latter is updated. Hence we have a dynamic structure which evolves as it is depicted by Figure 2.

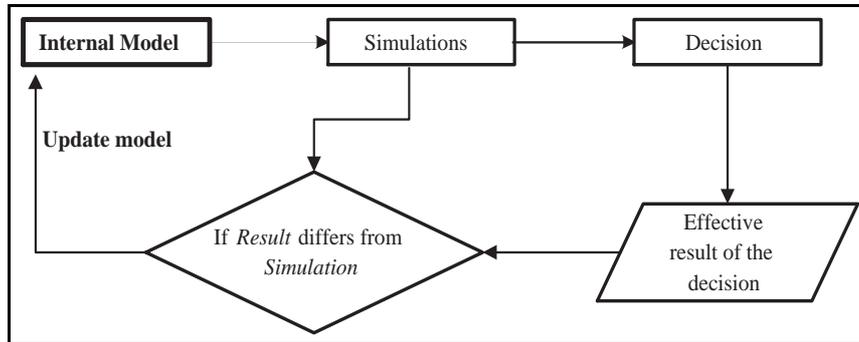


FIGURE 2. Dynamics of the internal model.

The representation of this internal model must therefore take into account the update of the model following the experience and the expectations of the firm. The main idea behind this approach is that “an intelligent being learns from experience, and then uses what it has learned to guide expectations in the future” (Dennett, 1998, p. 185) and, moreover, “. . . failed predictions can serve as well as overt reward as a basis for improvement” (Holland, 1998, p. 77). The idea we are trying to represent in our models is therefore a fairly obvious one: “an intelligent agent must engage in swift information-sensitive ‘planning’ which has the effect of producing reliable but not foolproof expectations of the effects of its actions” (Dennett, 1998, p. 193).

While this idea is quite obvious, its integration into models of R&D is problematic. This is the reason why purely adaptive models (see the preceding section) generally neglect this dynamic process of expectation formation. The representation of learning as the product of an evolutionary algorithm does permit the elaboration of better decision rules, but only by trial and error. The agent can only judge decisions which have been used before. On the contrary, the vision based on the dynamics of the internal model admits that agents can have a relatively precise (if not perfect) perception of the value of their decisions even if they have never been used before. This is made possible by means of simulations with the internal model. We must now attack the problem of the integration of this idea into economic models.

### 3.2. HOW TO REPRESENT THE EVOLUTION OF THE INTERNAL MODEL?

The standard way of formalizing such a model is given by subjective probabilities of Savage. In this case, the internal model of the firm corresponds to a set of conditional probability distributions. The update of this model can be imagined through successive least square estimations or Bayes rule.

The Bayesian approach has the advantage of not assuming any particular structure for the internal model. But it is very demanding in terms of agents’ rationality.

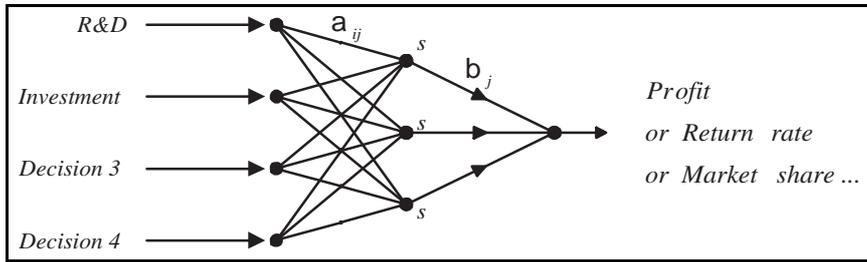


FIGURE 3. Feedforward ANN with one hidden layer.

Moreover, “there is substantial evidence that Bayes’ theorem lacks empirical relevance and hence its procedural justification is weak” (Salmon, 1995, p. 245).

Even if least squares can be used in simple representations of this internal model (see for example the use of this approach in Jonard and Yildizoglu, 1999 for the expectations about the return on physical investment), we need a tool as flexible as possible for the black-box approach. Ideally our representation should be independent of the structure and the parameterization of the internal model in order to incorporate only the most primitive dimensions of this model: its existence and its influence on the decisions of agents. An artificial neural network (ANN) is a good candidate to represent the dynamics of the internal model in a black-box approach. With only minimal structural assumptions, namely the list of dependant and explicative variables, and the structure of the hidden layer, it can represent the fact that the firm adjusts its internal model to the flow of experience. The results of the ANN are quite independent from the adopted activation function (only the speed of the training depends on it – see Masters (1993) for the discussion of properties of ANNs): for many practical problems even a very simple feed forward ANN with one hidden layer of few hidden nodes gives quite robust results.

An ANN provides a time varying flexible functional form that delivers an approximation of the connections between the inputs and the output of the internal model. This approximation is obtained by the calibration of the parameters of the ANN ( $a_{ij}$  and  $b_j$  in Fig. 3) according to the series of input and output data. These parameters reflect the intensity of the connections in the network. A better approximation can be achieved through the introduction of fictive nodes in the network (a hidden layer) which represent unobserved state variables, or more particularly, unobserved variables of the internal model of the agent. ANN thus covers a wide range of models from the simplest linear one when there is no hidden layers, to the increasingly sophisticated ones when the number of the hidden nodes increases. This number can even be used to represent the complexity of the agent’s internal model.

### 3.3. INTERNAL MODEL FOR R&D DECISIONS

The internal model used by firms for their R&D decisions could be represented by an ANN (see Fig. 4). Given a state of this ANN, the firm uses it to compare

alternative investment decisions in terms of their competitive outcome. On the basis of the results of the ANN, the agent takes a decision. This investment decision results in a certain performance determined by the dynamics of the industry. This decision and the corresponding outcome provides a new experiment for the firm. The use of this experiment for the update of the internal model corresponds to the learning.

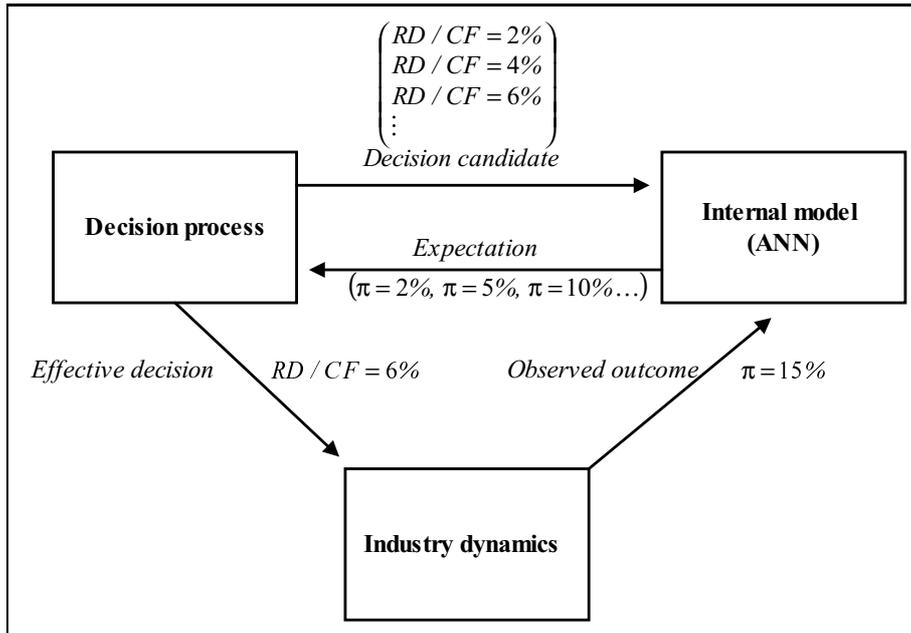


FIGURE 4. Expectations, decisions and learning.

At any point in time the state of the ANN is indeed given by parameters which represent the strength of connections. These parameters are calibrated at each period by using the past observations of inputs (decisions and indicators) and outputs (indicators of competitiveness or profitability). Then the agent provides this ANN with different investment hypothesis and compares the resulting outputs:

$$\text{Input : } "(RD/CF = x\%)" \longrightarrow \text{Output : } "\pi = y\%".$$

where  $RD/CF$  represents the R&D- cash flow ratio and  $\pi$ , the profit rate of the firm. This expectation ( $x\% \rightarrow y\%$ ) can be compared with other hypothesis in order to make a decision through the relevant decision process. With bounded rationality, the firm will choose the first hypothesis with a *satisfying* outcome. This decision and its effective result (the observed profit rate) can again be used to update the internal model (to calibrate the ANN).

As input data one could use actual or lagged values of different relevant variables: R&D investments (cumulated or not), capital stock (size), market share,

known demand characteristics, the degree of maturity of the technology. Cumulated or lagged values of R&D investments can be particularly relevant since the innovation process is cumulative and R&D does not have immediate reward. As to competitiveness we can consider different indicators such as relative mark-up, variation rate of the productivity or the variation of the profit rate.

#### 4. A MODEL OF ADAPTIVE LEARNING AND EXPECTATIONS

In order to test the validity of the approach proposed in the preceding section, I will now present an industry model where firms elaborate their profit-rate expectations from their individual internal model. The presence of such a process is modeled under the form of a firm-specific artificial neural network. As in Yildizoglu (2001), the exploration by firms of the R&D-strategy space is represented through a genetic algorithm.

##### 4.1. A SIMPLE MODEL

This model is very close to Yildizoglu (2001) that includes *learning* firms in a simplified version of the Nelson and Winter (1982) model. As in the Nelson & Winter model, this model concerns an industry producing a homogeneous good and facing a decreasing market demand. The only production factor is the physical capital and the technology has constant returns to scale (it is linear). In each period, each firm dedicates a fraction of its gross profits to the R&D activity and the residual of the profits are used for the expansion of the physical capital. The R&D investment of the firm is necessary for the imitative and innovative activity of the firm, and these are the only source of productivity gains in the model.

In the previous version of the model, learning firms (**GenFirms**) that use a Genetic Algorithm for exploring the strategy space (the R&D/Profits ratios) are introduced. In order to assess the fitness of each decision rule (R&D rates) these firms effectively use them during a number of periods (learning period), in order to take into account the delayed and random results of the innovation process. The results of the model show that even with such a crude exploration strategy, the presence learning firms is a source of technological performance for the industry and of a competitive edge for themselves.

Of course, the assumption of a configuration where firms effectively use any idea that they can happen to have on their strategy is very far from realistic. As we have proposed in the preceding sections, they will rather have an internal model that will hint them about the possible value of each strategy, even before the first use of it. The presence of such a model on the base of which the expectations are formulated by the firm is included in the actual model, using a firm-specific ANN (see Masters, 1993; Watson, 1997) for each firm. **NNGenFirms** use such a procedure for learning.

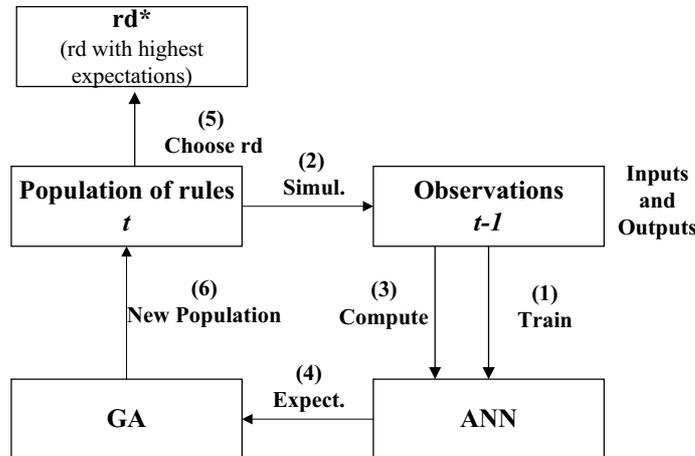


FIGURE 5. GA and ANN based learning.

Figure 5 summarizes the workings of the learning block of the **NNGenFirms** (a detailed description of the model in the standard SUN JavaDoc format can be read from my web site<sup>2</sup>).

Expectations are based on an internal model that connects variables observable by the firm (Inputs vector of the ANN) to the desired output. In each period:

1. past observations of inputs and the output are used to train the network (adapting the internal model to the observations of the firm - arrow 1);
2. using the freshly trained ANN, the firm simulates the possible outcome of each possible R&D rate in the actual strategy population, combining this rate with the actual values of other inputs: for each rule in the population, an inputs vector composed of this value of R&D rate and last observations on other inputs are provided to the ANN (arrow 2). The ANN uses this input to compute the corresponding value of the output given its actual state (arrow 3) and the GA uses this value as the expected fitness of the rule (arrow 4);
3. the R&D strategy with the highest expected profit rate is used by the firm to determine the amount of its R&D investment (arrow 5);
4. the GA uses the expected fitness of the rules to generate a new population through selection-crossover-mutation (arrow 6);
5. the R&D investment of the firm determines the stochastic outcomes of the imitation and the innovation: if successful, these activities can provide the firm with a higher productivity of the capital (the only production factor);
6. the residual gross profits (profits – R&D investment) determines the investment of the firm in physical capital and, hence, the capital stock of the next period;

<sup>2</sup><http://beagle.montesquieu.u-bordeaux.fr/yildi/learnindnn/doc/>

7. the simulation is carried on in the next period.

The **NWFirms** use a fixed R&D rate in stages 5–7 in a way similar to the original Nelson and Winter Model. The **GenFirms** use the same R&D rate during the learning period and, once all rules in the population are used, they proceed with stage 4.

For the NNGenFirms, the GA and the ANN are hence complementary: the GA represents the adaptive explorations of the firm and the ANN provides the GA with the expectations that are necessary for the *ex ante* evaluation of strategies. This model complies with our announced objective: connecting adaptive behaviour with expectations.

#### 4.2. RESULTS OF THE EXPERIMENTS

Since we aim to derive results independent from a particular sequence of random numbers, a batch of 20 simulations, of 2000 periods each, is run for each configuration of the model. Observations have been saved every 40 periods. Each possible configuration of the industry is hence represented by a sample of 1000 observations. The relevant dimensions (*e.g.* technical progress, concentration) of resulting samples are compared by way of non-parametric tests (the non-parametric Mann–Whitney test, see for instance ch.18 in Watson *et al.*, 1993). For convenience, results are presented as **box plots** where the box gives the central 50% of the sample centered around the median: the box hence gives the first, second and third quartiles ( $Q_1, Q_2, Q_3$ ) of the distribution. The whiskers give the significant minimum and the significant maximum of the distribution. Each box contains the whole history of the industry for all simulations for each corresponding configuration.

This protocol allows the qualitative comparison of different industry configurations. Different indicators are used for these comparisons.

Quite standard indicators are used for the comparison of performance of industries:

- welfare indicators: market price and average gross profits;
- technical efficiency indicators: average productivity and maximal productivity.

I compare three industry configuration:

**nwpure:** An industry composed of 30 NWFirms investing each 10% of gross profits in R&D.

**genpure:** An industry composed of 30 GenFirms, each using a GA with 12 chromosomes of 8 genes. Each rule is evaluated during 3 periods (the learning period).

**ngenpure:** An industry composed of 30 NNGenFirms that use the same GA as the GenFirms but use an ANN with 4 hidden nodes for evaluating each rule. The inputs of the ANN are:

- the R&D rate;

- a moving average of the last 5 R&D rates, taking into account the delayed effect of the R&D;
- the variation rate of the physical capital of the firm ( $\Delta K/K$ ), taking into account the impact of the capital expansion on profit;
- $\max\{0, -\Delta p/p\}$  that takes into account the impact of the price ( $p$ ) on the profit rate (normalized in order to give values in accordance with other inputs);

Only the data from last 10 periods is used for training the ANN in each period.

Figure 6 summarizes the results concerning the market performance of these industries. Figure (b) shows that learning is costlier for the NNGenFirms because of the imprecision of the ANN during the beginning of the history. They can hence suffer from higher losses than the GenFirms. But they can also attain much higher profits than the GenFirms (that perform better than the NWFirms as I show in Yildizoglu (2001)). The Mann-Whitney tests gives the following ordering for the medians of the average profits in these industries:  $nwpure < genpure < ngenpure$ . Globally, we observe highest average profits in industries composed of NNGenFirms. Again, the comparison of the distributions of the market price (Fig. 6a) is clarified using Mann-Whitney tests: the lowest prices are observed in industries composed of NNGenFirms. The social welfare is consequently the highest in such industries in comparison with other configurations. There is not a contradiction between these figures (lowest prices but higher profits) if the productivities are higher in *ngenpure* industries because this would imply lower unit costs.

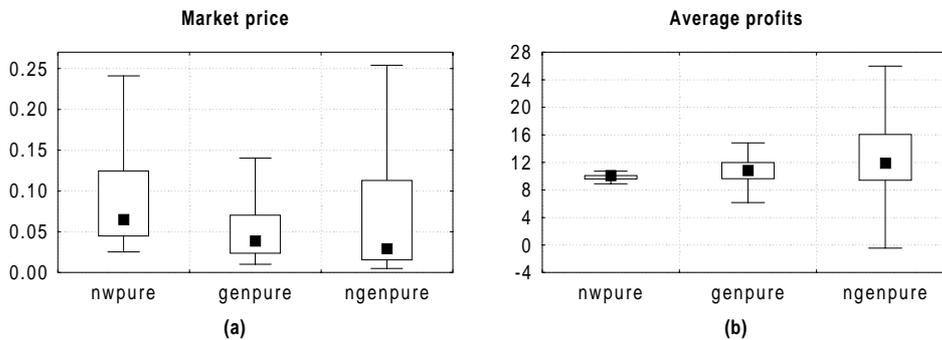


FIGURE 6. The relative market performance of industries with NNGenFirm.

Figure 7a confirms this result: *ngenpure* industries attain higher maximal profits. Figure 7b shows that *nwpure* industries are not very efficient in R&D and this mainly results from constant expansion of the physical capital even if this tremendously lowers the price and hence, the productivities rate. This figure also indicates an R&D process quite inefficient for the GenFirms: they do not spend enough on R&D to attain the same maximal productivities as the NNGenFirms.

The introduction of the expectations enhances the performance of the adaptive process represented by the GA.

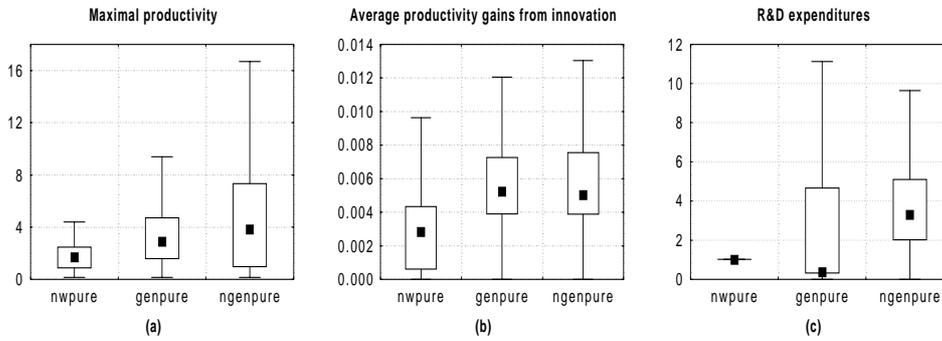


FIGURE 7. The relative technological performance of NGenFirms.

It would be interesting to test the sensitivity of these results to the structure of the ANN. Figure 8 represents the performances of *ngenpure* industries with different ANNs (aggregated over the batch of 20 simulations of 2000 periods for each configuration – *nhi* corresponds to ANNs with *i* hidden nodes). These results show that the performance of the *ngenpure* industries does not strongly depend on the structure of the ANN: the structure with 4 hidden nodes corresponds to slightly (but in a statistically significant way) better performances but we do not observe a regular relation between the number of hidden nodes and the market and technological performance of NGenFirms. As a consequence, these results are quite robust in respect with the structure of the ANN (especially when completed with the insensitivity to the structure of the activation function, as we have earlier discussed): the distributions of maximal productivities for the industries *ngenpure* in Figure 7a and *nh4* in Figure 8a are very similar even if they result from different sets of simulations.

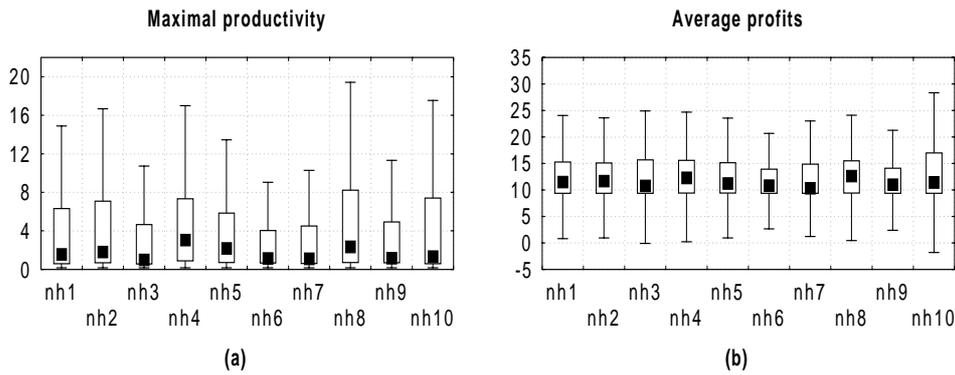


FIGURE 8. Sensitivity to the variation of the number of hidden nodes.

## 5. CONCLUSIONS

This article establishes the necessity and the possibility of complementing adaptive learning mechanisms with adaptive expectations in the modeling of R&D strategies of firms with bounded rationality.

We start from a trivial observation: if firms get engaged in a costly and rather inflexible innovation process, that is because they expect to get something back from it. These expectations and their appreciation of the technological trajectory of their industries will strongly condition the importance of the R&D expenditures in their budget. This observation does naturally not imply that we should assume the simplest of the worlds and model these expectations with rational expectations: Firms are generally very far from having a crystal-clear picture of the dynamics of the economy, but they at least have some picture of the economy, and it guides their decisions. Even the adaptive decision processes of boundedly rational firms are guided by expectations.

This article proposes that we should model the presence of this picture and its evolution as a consequence of the observations of the firm using a black box approach based on the concept of *internal model*, and that the presence (but not the structure) of this internal model can be represented by a firm specific artificial neural network. The second part of the article develops such a model based on Yildizoglu (2001): the adaptive decisions of the firms are represented by a genetic algorithm as in the preceding article but the fitness of each strategy is *ex ante* appreciated by the firm using *as if* experiments through a firm specific artificial neural network. I compare the performance of industries composed of such firms with the performances of industries with firms using fixed rules *à la* Nelson & Winter and others composed of firms that effectively use each rule in order to appreciate its value. The results clearly show that we observe higher welfare and technological performance if we take into account the expectations of the firms. Moreover, these results are very robust and they do not depend on the structure of the ANN (the number of hidden nodes). These results confirm our vision concerning the necessity of taking into account the fact that firms form expectations about their environment and these expectations are a source of better performance. If we ignore them, we can underestimate in our models the performance of the firms. These are only first results with this model. A more in depth analysis of the dynamics of the model should reveal even richer properties of learning firms.

*Acknowledgements.* I would like to thank Vanessa Oltra, Bernard Ancori, Ragip Ege, Gerald Silverberg and two anonymous referees for stimulating comments on a previous conceptual version of this work. The support of CNRS research program “Les enjeux économiques de l’innovation” is gratefully acknowledged.

## REFERENCES

- Ballot G. and Taymaz E. (1997), The dynamics of firms in micro to macro model with training. *J. Evol. Econom.* **4**.
- Ballot G. and Taymaz E. (1999) Technological change, learning and macroeconomic coordination: An evolutionary model. *J. Artificial Societies and Social Simulation* **2**, <http://www.soc.surrey.ac.uk/JASS/2/2/3.html>

- Dennett D.C. (1998) *Brainchildren. Essays on Designing Minds*. Penguin Books, London.
- Dosi G, Faggiolo. G. and Marengo L. (1999) Learning in evolutionary environments, in K. Dopfer, ed., *Evolutionary Foundations of Economics*. Cambridge University Press, Cambridge.
- Goldberg D.E. (1991) *Genetic Algorithms*. Addison-Wesley, Reading: MA.
- Hodgson G.M. (1994) *Economics and Evolution*. Polity Press, Cambridge: UK.
- Hofstadter D.R. (1979) *Gödel, Escher, Bach: An Eternal Golden Braid*. Basic Books, New York.
- Holland J.H. (1996) *Hidden Order. How Adaptation Builds Complexity*. Addison-Wesley, Reading (MA).
- Holland J.H. (1998) *Emergence. From Chaos to Order*. Addison-Wesley, Reading (MA).
- Jonard N. and Yildizoglu M. (1999) Sources of technological diversity. Cahiers de l'Innovation 99030, CNRS, Paris.
- Knight F.H. (1921) *Risk, Uncertainty and Profits*. Number Reprint, Chicago University Press, Chicago.
- Kwasnicki W. (1998) Skewed distributions of firm sizes - an evolutionary perspective. *Structural Change and Economic Dynamics* **9**(1), 135–158.
- Kwasnicki W. and Kwasnicka H. (1992) Market, innovation, competition. An evolutionary model of industrial dynamics. *J. Economic Behavior and Organization* **19**, 343–368.
- Masters T. (1993) *Practical Neural Network Recipes in C++*. Academic Press, San Diego: CA.
- Mitchell M. (1996) *An Introduction to Genetic Algorithms*. The MIT Press, Cambridge: MA.
- Nelson R.R. and Winter S. (1982) *An Evolutionary Theory of Economic Change*. The Belknap Press of Harvard University, London.
- Rubinstein A. (1998) *Modeling Bounded Rationality*. The MIT Press, Cambridge: MA.
- Salmon M. (1995) Bounded rationality and learning: Procedural learning, in Kirman, Alan; Salmon, Mark, eds., *Learning and Rationality in Economics*. Blackwell, Oxford, pp. 236–275.
- Sargent T.J. (1995) *Bounded Rationality in Macroeconomics*. Oxford University Press.
- Silverberg G. and Verspagen B. (1995) Long term cyclical variations of catching up and falling behind. An evolutionary model. *The J. Evol. Econom.* **5**, 209–22.
- Silverberg G. and Verspagen B. (1999) Evolutionary theorizing on economic growth, in K. Dopfer, ed., *The Evolutionary Principles of Economics*. Cambridge University Press, Norwell, MA.
- Silverberg G. and Verspagen B. (1996) From the artificial to the endogenous, in M. Helmstadter, Ernst; Perlman, ed., *Behavioral norms, technological progress, and economic dynamics: Studies in Schumpeterian economics*. University of Michigan Press, Ann Arbor.
- Simon H.A. (1958) The role of expectations in adaptive or behavioristic model, in M. Bowman, ed., *Expectations, Uncertainty and Business Behavior*. Social Science Council, New York, pp. 49–58.
- Simon H.A. (1976) From substantial to procedural rationality, in S.J. Latsis, ed., *Method and Appraisal in Economics*. Cambridge University Press, Cambridge, pp. 129–148.
- Simon H.A. (1982) *Models of Bounded Rationality*, Vol. 2. Behavioral Economics and Business Organization, The MIT Press, Cambridge: MA.
- Watson C.J., Billingsley D.J., Croft D.J. and Huntsberger D.V. (1993) *Statistics for Management and Economics*. Fifth edition, Allyn and Bacon, Boston.
- Watson M. (1997) *Intelligent Java Applications*. Morgan Kaufmann, San Francisco: CA.
- Wilson S.W. (1995) Classifier Fitness Based on Accuracy. *Evolutionary Computation* **3**(2), 149–175. <http://prediction-dynamics.com/>
- Winter S. (1984) Schumpeterian competition in alternative technological regimes. *J. Economic Behavior and Organization* **5**, 287–320.
- Yildizoglu, M. (2001) Competing R&D strategies in an evolutionary industry model. *Forthcoming in Computational Economics*. Available at <http://yildizoglu.montesquieu.u-bordeaux.fr/>