

Learning and Expectations in R&D Decisions

Vanessa Oltra[†] and Murat Yıldızoğlu^{‡1}

[†]IFREDE-E3I.

Université Montesquieu Bordeaux IV
Aven. Leon Duguit, 33608 PESSAC.

[‡]BETA-THEME (UMR 7522 - CNRS).

PEGE

Université Louis Pasteur.
63, Avenue de la Forêt-Noire
67085 Strasbourg Cedex.
e-mail: yildi@cournot.u-strasbg.fr

¹We would like to thank Bernard Ancori, Ragip Ege and Gerald Silverberg for stimulating comments. The support of CNRS research program “Les enjeux économiques de l’innovation” is gratefully acknowledged.

Abstract

In this methodological work we 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, we propose a *black box* strategy through the concept of “internal model”. The last part of the article uses the artificial neural networks as a tool to model the presence and the evolution of the internal model of agents.

Keywords: Bounded rationality, learning, expectations, innovation dynamics.

And one of the deepest, one of the most general functions of living organisms is to look ahead, to produce future as Paul Valéry put 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 concerning 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 corresponds to a decision which is turned to the future and, as a consequence, which involves an expectation dimension.

In the meantime, 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 clarify the role of uncertainty in R&D decisions, we should distinguish between two types of uncertainty:

- 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;

- 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, we 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). Finally the combination of expectation with adaptation in the presence of uncertainty appears as relevant for any types of innovation (incremental or radical).

We propose to explore the modelling of the determination of the level of R&D investment of firms. This means that we are not going to tackle the decision of being an innovator or not, nor the adoption of a new technology. We are going to 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. Our interest is to analyze how expectation and adaptation can be combined in the modelling of R&D investment rules. In the literature both dimensions are generally split up: rational expectations are assumed in neo-classical models whereas alternative approaches (institutional and/or evolutionary) generally adopt a purely adaptive representation.

In the next section, we discuss rational expectation models. We consider two different models of innovation: Aghion and Howitt (1992) and Ericsson and Pakes (1995). The former is a canonical model of endogenous growth while the latter is a very sophisticated model of industry dynamics. Both endogenize innovation as a stochastic outcome of R&D investment of firms. We show that this class of models deals with risk (and not with uncertainty) and does not cope with the formation and the adaptation of expectations. In section 3, we focus on models based on adaptive decision rules. We show that recent models of innovation developed in the evolutionary framework tend to overlook the role of expectation in innovation decisions. In this class of models, the decision rules are purely adaptive, and as a consequence do not take into account the impact upon investment of the expectations of firms concerning the relationships between R&D and competitiveness. Section 4 is devoted to the modelling trade-off between expectation and adaptation. We 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.

2 Equilibrium and risk: rational expectation models

Given our interest in the sources of uncertainty, we focus on non-deterministic models of innovation. Our purpose is to study the decision rules governing R&D investment in the models in which innovation is endogenized as a stochastic process. Two types of models come within this perspective: endogenous growth models and industrial dynamics models. For our analysis of the standard approach of R&D decisions, we choose the canonical model of endogenous growth of Aghion and Howitt (1992) and the model of industrial dynamics developed by Ericsson and Pakes (1995). These two models illustrate the modelling of R&D decisions in an equilibrium approach.

2.1 Macroeconomic dynamics and rational expectations

Aghion and Howitt (1992) present a model of endogenous growth that is based on Shumpeter's idea of creative destruction. The particular process it focuses on is research activities aimed at producing innovations which consist in developing better products. In this model, the inputs of research are skilled labor and specialized labor. A successful research firm obtains a patent on its innovation which it sells to an intermediate firm that becomes a monopolist until the next innovation occurs. At any point in time the main decision is how to allocate the fixed

flow of skilled labor between manufacturing and research.

Labor used in research produces a random sequence of innovations that follows a Poisson law. The Poisson arrival rate of innovations in the economy at any period is a function of the flow of skilled labor used in research. The probability law of innovations is stationary and research firms know the Poisson arrival rate. Moreover each innovation increases the productivity by a constant factor. Thus the only uncertainty for research firms concerns the time of innovation. Given that firms know the Poisson arrival rate, they do not have to cope with uncertainty but with technological risk.

Concerning the market environment, the occurrence of innovation for each firm is independent of the inputs of other research firms. An innovative firm obtains systematically a patent which it sells to an intermediate firm. The patent price is equal to the expected present value of the monopoly rents of the intermediate firm. Given that each innovation destroys the monopoly rents resulting from the previous innovation, the only uncertainty concerns the length of the interval of monopoly rents. Thus the value of an innovation is the expected present value of the flow of monopoly profits generated by the innovation over an interval whose length is exponentially distributed. Firms know the function which determines the parameter of this exponential distribution. This knowledge enables them to maximize the expected present value of profit.

But to calculate this expected value firms have to take into account the cost of inputs which is the wage rate of labor. More precisely the expected wage of period $t + 1$ is used to calculate the expected value of the $(t + 1)$ th innovation. This means that firms must be able to forecast the evolution of the wage rate. In order to ensure intertemporal coordination the authors assume that agents have perfect foresights. This assumption is a necessary condition for the intertemporal equilibrium. Thus firms know the probability laws of the stochastic processes and perfectly forecast the evolution of the cost of inputs. It implies that firms perfectly know the model of the economy. As a consequence they do not cope with uncertainty, but make optimal decisions under risk. In this context, there is no need for any adaptive dimension. Firms do not have to adapt to unpredicted changes since they perfectly predict future outcomes. Moreover there is no memory in the technology since the arrival rate of innovations depends only on the current flow of input to research. It means that there is no accumulation nor irreversibilities in the research process of firms. R&D decisions are totally turned to the future which is predicted and described by stationary probability distributions.

2.2 Industrial dynamics and rational expectations

Ericson and Pakes (1995) formalize research activities in an other framework by using Markov processes. Their model of industry dynamics is based upon a stochastic model of the entry and growth of firms that invest in research and exploration activities to enhance their capability to earn profits. The stochastic outcome of the firm's investment, the success of other firms in the industry and the competitive pressure from outside the industry determine the profitability of each firm.

Technological opportunities provided by the industry are open to all, so that the only distinction among firms is their efficiency in exploiting these opportunities. The efficiency of each firm is measured relatively to the efficiency of the other firms in the industry and to the competition from outside the industry. The efficiency of firms and the industry structure evolve as a result of the outcomes of firms' investment and with changes in the market environment in which it is embedded. The model includes an exogenous stochastic process which reflects the improvements made by competition outside the industry (linked to the evolution of demand,

input costs, science and technology). At each period the state of the industry is given by the efficiency of firms and the market structure which is described in terms of the number of firms at each state of efficiency.

The level of investment of firms is chosen so as to maximize the expected present discounted value of profits as a function of all information available to firms. The authors assume that this information includes:

- the history of all past states of the industry,
- the history of the firm's own past investment decisions,
- the current state of the industry,
- the probability laws governing the evolution of that state over time, including the law governing the impact of the firm's own investment on that evolution.

In this model technological uncertainty concerns the stochastic outcomes of firms' research investment and there are two sources of market uncertainty which are the evolution of competition outside the industry and the process of entry of new firms in the industry. Since firms know the probability laws governing these factors, they can maximize the expected future profits generated by investment. Thus in spite of the complexity of the model, firms are sufficiently informed to cope optimally with risk. This property is due to non-trivial behavioral assumptions which seem to us difficult to justify empirically and of poor relevance for R&D decisions.

Moreover in order to determine their optimal investment strategy, firms need to form expectations about the evolution of the state of efficiency of the other firms in or entering the industry. The distribution used to form these expectations is derived from the firm's perception of the Markov process governing the state of efficiency of its competitors. It corresponds to the firm's beliefs about how the structure of the industry will change. The industry is said to be in dynamic equilibrium when the process generating the change in industry structure is accurately reflected in the beliefs of each firm. The authors refer to a rational expectations equilibrium where optimal decisions are based on the true distribution of future states generated by the optimal behavior of all firms.

The assumption of rational expectation insures intertemporal coordination of the future decisions of firms. It also means that the firms know perfectly the model of the economy or that the model they use to form their expectations does not lead to systematic errors in prediction. The latter argument implies that errors in prediction are not correlated and that the distribution of errors is centered around a mean equal to zero. Thus even if errors in prediction occur, the agents do not change the model they use to form their expectations. Under the rational expectation assumption, expectations lead agents to make decisions whose outcomes do not refute their predictions (Gaffard (1997)). In the model of Ericson and Pakes, it means that in spite of the complexity of interactions among firms and the risk associated to research investment, firms have the ability to predict the evolution of the state of the industry by using a "good" model of the economy. This model is not adapted through time since it does not entail systematic errors in prediction.

The presentation of Aghion and Howitt (1992) and Ericson and Pakes (1995) shows that none of these models deals with actual technological or market uncertainty. Even if the authors introduce stochastic processes to represent innovations, the assumption according to which firms know the probability laws, that are generally stationary, implies that firms deal with risk and not with uncertainty. We usually find this assumption in the neoclassical approach to innovation. In this framework, knowledge about probability distributions is necessary to firms in order to be able to maximize the expected present value of profits.

We also emphasize that the concept of equilibrium with rational expectations or perfect foresights, 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 which are guided by the agents' vision of the technological change process.

3 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 centre 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 when 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.

3.1 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 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 (1995, 1995b) 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.

<i>Population at date t</i>	<i>Fitness: $f(x)=x^2$</i>	<i>Expected number: $f(x)/Mean_f$</i>	<i>Effective number proportionally drawn</i>	<i>New population</i>
1) 00011 (=3)	1) 9	1) 0.0	1) 0	1) 1 0111 (=23)
2) 01100 (=12)	2) 144	2) 0.6	2) 1	2) <u>011 00</u> (=12)
3) 10111 (=23)	3) 529	3) 2.4	3) 2	3) 10111 (=23)
	Mean $f_t=227$	Sum=3	Sum=3	
<i>Crossover (3)-(2) at bit 3</i>		<i>Mutation (1)</i>		<i>Population at date t+1</i>
1) 1 0111 (=23)		1) 1 1111 (=31 *)		1) 1 1111 (=31)
2) <u>011 11</u> (=15)		2) 011 11 (=15)		2) 011 11 (=15)
3) 10100 (=20)		3) 10100 (=20)		3) 10100 (=20)

Simple application of GAs to optimisation of the function $f(x)=x^2$ 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->1 or 1->0. **This process is controlled by: population size, bit-string size, probability of crossover and probability of mutation.**

Figure 1: A simple example of genetic algorithm

3.2 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) and Kwasnicki (1998)) use a representation of the learning of firms already very close to genetic algorithms (see Figure 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 modelled 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, pp.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 & 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. 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) and Mitchell (1996)). 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 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”.

4 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 resign 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) and Dennett (1998)).

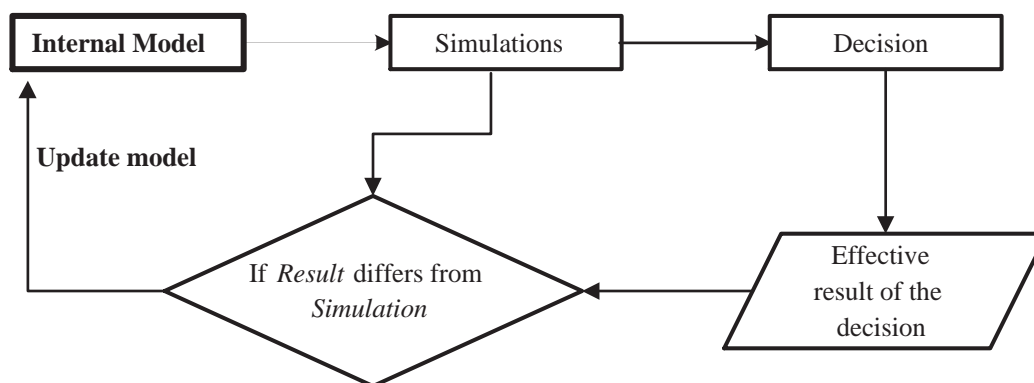


Figure 2: Dynamics of the internal model

4.1 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 an 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.

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 our 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.

4.2 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.

With least square estimations, we must have an idea of the structure of the model of the firm and assume that the agent does not modify this structure in the updating process, but only its quantitative elements (mainly the coefficients of the equation). Even if this is very restrictive, we will consider some possible applications of this methodology to R&D

decisions. Two problems limit the relevance of this process: the decreasing importance through time of more recent experiences, and the impossibility for the agent to modify the structure of the internal model as a result of its experiments (Salmon (1995), Sargent (1995)). The first limitation can easily be neutralized by assuming a fixed memory horizon for the agent (estimations over last n periods instead of the whole history). The second limitation cannot be handled by least squares and so limits the relevance of this approach to relatively stable environments (characterized, for example, by incremental technical progress).

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




Figure 3: Feedforward ANN with one hidden layer

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 (1998) 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, it can represent the fact that the firm adjusts its internal model to the flow of experience.

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 Figure 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.

We can now propose two examples of application of the black-box approach to R&D decisions.

4.3 Internal model for R&D decisions

Our first example proposes a very simple model of R&D decisions where the internal model is represented by a linear model. This model is adapted to the flow of experience through incremental least squares. The second example is dedicated to the possible use of ANN in a very simple framework.

In order to represent the expectations of the firm concerning the R&D-profitability connection, we can use incremental least square estimation of a linear model. In that case the update of the internal model is based on a constant structure. The only elements which are really updated according to experience are the coefficients of the model.

Expectations of the agent are based on the predictions of the estimated model. In a direct use of such a model the agent can estimate the return on R&D investments and compare it to the return of other investment possibilities (physical capital, financial). But R&D has not necessarily an instantaneous return and this approach neglects the delayed impact of R&D on profitability or competitiveness. Thus an under-estimation of the return on R&D can be observed. This can be overcome by the use of the cumulated R&D investment as an explicative variable. The linear structure of the model makes useless the test of different hypothesis (amounts) of R&D investment. The impact of R&D on competitiveness is given by the same coefficients for all values of this investment. Nevertheless, these limitations and the one which comes from the constant structure of the model should not eliminate this modelling possibility, especially in the case of incremental innovation.

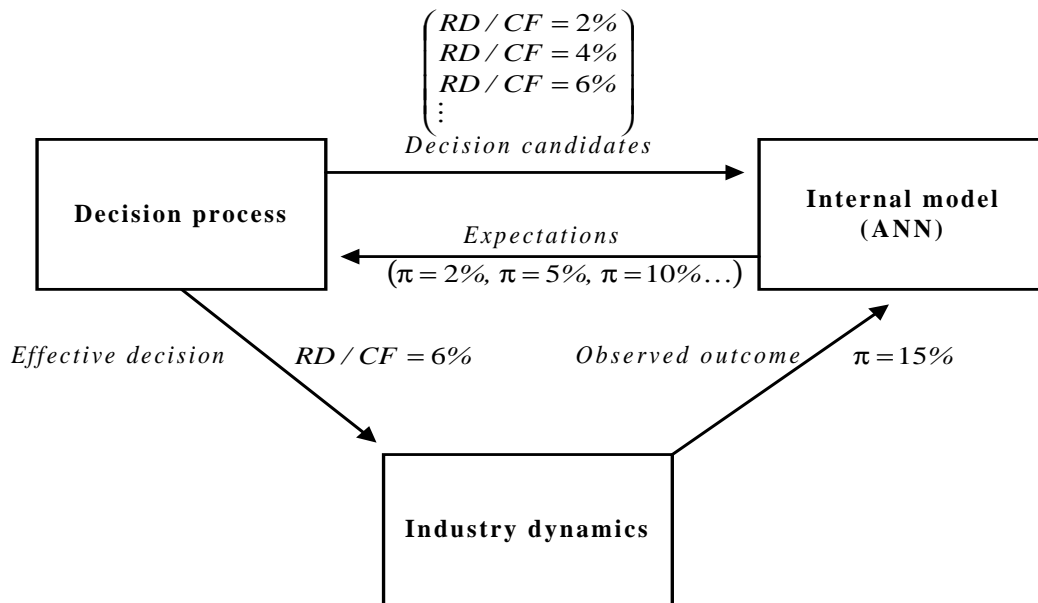


Figure 4: Expectations, decisions and learning

For a more general approach, the internal model used by firms for their R&D decisions could be represented by an ANN (see Figure 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.

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:

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

where RD/CF represents the R&D- cash flow ratio and π , 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. Because of bounded rationality, the firm content itself with testing a limited set of hypothesis and cannot generally attain a global optimum.

5 Conclusion

As a conclusion we would like to stress that we can leave a room for expectations in models of R&D without assuming that firms know the exact model of the economy. The important point is that firms form expectations by using their internal model and adapt this model according to their experience and observations. This approach opens up new perspectives to cope with the trade-off between expectations and adaptation in R&D models.

The approach we propose in this article is also close to the workings of a classifier system but it has the advantage of explicitly modeling the formation of expectations. This separation of expectations and decisions also allows for the possibility of different frameworks for the modeling of these two processes and should enrich our analysis of learning processes.

References

- Aghion, P. and Howitt, P. (1992) A Model of Growth Through Creative Destruction. *Econometrica*, 60, 323-351.
- Arrow, J.K. (1951). Alternative Approaches to the Theory of Choice in Risk-Taking Situations. *Econometrica*, 19, 404-37.
- Ballot, G. and Taymaz, E. (1997) The Dynamics of Firms in Micro to Macro Model with Training. *Journal of Evolutionary Economics*, 4 .
- Ballot, G., and Taymaz, E. (1999) Technological Change, Learning and Macroeconomic Coordination: An Evolutionary Model. *Journal of Artificial Societies and Social Simulation*, <http://www.soc.surrey.ac.uk/JASS/2/2/3.html>,, 2.
- Dennett, D.C. (1998) *Brainchildren. Essays on Designing Minds*. London, Penguin Books.

- Dosi, et al. (1999) Learning in Evolutionary Environments. In K. Dopfer (ed.), *Evolutionary Foundations of Economics*. Cambridge, Cambridge University Press.
- Ericsson, R. and Pakes, A. (1995). Markov-Perfect Industry Dynamics: A Framework for Empirical Work. *Review of Economic Studies*, 62, 53-82 .
- Gaffard, J. (1997) *Croissance et fluctuations économiques*. Paris, Montchrestien.
- Goldberg, D.E. (1991) *Genetic Algorithms*. Reading: MA, Addison-Wesley.
- Hodgson, G.M. (1994) *Economics and Evolution*. Cambridge: UK, Polity Press.
- Hofstadter, D.R. (1979) *Gödel, Escher, Bach: An Eternal Golden Braid*. New York, Basic Books.
- Holland, J.H. (1996) *Hidden Order. How Adaptation Builds Complexity*. New York, Helix Books (Addison-Wesley).
- Holland, J.H. (1998) *Emergence. From Chaos to Order*. Reading: MA, Addison-Wesley.
- Jonard, N. and Yildizoglu, M. (1998) Sources of Technological Diversity, Strasbourg, mimeo, BETA.
- Knight, F.H. (1921) *Risk, Uncertainty and Profits*. Chicago, Chicago University Press.
- 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. *Journal of Economic Behavior and Organization*, 19, 343-368
- Loasby, B.J. (1986) Organisation, competition, and the growth of knowledge. In R.N. Langlois (ed.), *Economics as a Process*. Cambridge, Cambridge University Press.
- Marengo, L. (1992) *Knowledge, Coordination and Learning in an Adaptive Model of the Firm*, DPhil Thesis, SPRU, Sussex University.
- Mitchell, M. (1996) *An Introduction to Genetic Algorithms*. Cambridge: MA : The MIT Press.
- Nelson, R.R. and Winter, S. (1982). *An Evolutionary Theory of Economic Change*. London: The Belknap Press of Harvard University.
- Rubinstein, A. (1998) *Modeling Bounded Rationality*. Cambridge: MA, The MIT Press.
- Salmon, M. (1995) Bounded Rationality and Learning: Procedural Learning. In A. Kirman and M. Salmon (eds.), *Learning and Rationality in Economics* (pp.. 236-275). Oxford, Blackwell.
- Sargent, T.J. (1995) *Bounded Rationality in Macroeconomics*, Oxford University Press .
- Silverberg, G., and Verspagen, B. (1995) *An Evolutionary Model of Long Term Cyclical Variations of Catching Up and Falling Behind*, IIASA .
- Silverberg, G. and Verspagen, B. (1995b) From the Artificial to the Endogenous: Modelling Evolutionary Adaptation and Economic Growth. *IIASA Working Paper*, 95-08, February.
- Silverberg, G., and Verspagen, B. (1999) Evolutionary Theorizing on Economic Growth. In K. Dopfer (ed.), *The Evolutionary Principles of Economics*. Norwell: MA, Cambridge University Press.
- Simon, H.A. (1958) The Role of Expectations in Adaptive or Behavioristic Model. In M.J. Bowman (ed.), *Expectations, Uncertainty and Business Behavior* (pp.. 49-58). New York, Social Science Council.
- Simon, H.A. (1976) From Substantial to Procedural Rationality. In S.J. Latsis (ed.), *Method and Appraisal in Economics* (pp.. 129-148). Cambridge, Cambridge University Press.
- Simon, H.A. (1982) *Models of Bounded Rationality*. Cambridge: MA : The MIT Press.
- Winter, S. (1984) Schumpeterian Competition in Alternative Technological Regimes. *Journal of Economic Behavior and Organization*, 5, 287-320 .