Competing R&D Strategies in an Evolutionary Industry Model

Murat Yıldızoğlu[1]
BETA-THEME (UMR 7522 of CNRS)
PEGE –Université Louis Pasteur
61, Avenue de la Forêt Noire. F-67000.
Fax. 333 90 41 40 50
e-mail: yildi@cournot.u-strasbg.fr

September 11, 1999 v.2.1

[1] I would like to thank Nicolas Jonard, Patrick Llerena and three anonymous referees for stimulating comments. Support for this research by the CNRS Research Program “Les enjeux économiques de l’innovation” is gratefully acknowledged.

Abstract

This article aims to test the relevance of learning through Genetic Algorithms, in opposition with fixed R&D rules, in a simplified version of the evolutionary industry model of Nelson and Winter. These two R&D strategies are compared from the points of view of industry performance (welfare) and firms' relative performance (competitive edge): the results of simulations clearly show that learning is a source of technological and social efficiency as well as a mean for market domination.
Keywords: Innovation, Industry dynamics, Bounded rationality, Learning, Genetic algorithms

1 Introduction

Research and development (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 with the 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 (see Oltra & Yildizoglu [1998]). Modelling R&D decisions must consequently rely on decision rules more sophisticated than the fixed rules that have been traditionally adopted in models of technology dynamics under bounded rationality assumption (see, for example, the models in Nelson & Winter [1982]).

As a matter of fact, fixed rules are not a necessary characteristic of decisions under bounded rationality: the main characteristic is procedural rationality (Simon [1976]). Hence, bounded rationality does not preclude the tendency of agents to adjust their decisions to the evolution of their (technological and competitive) environment. Even if they do not search for the globally best solutions, agents learn from their experience and this learning allows them to *fine-tune* their decisions. Consequently, one must search for a better way of modelling decisions in order to take into account this individual learning.

Of course, one could choose to implement a simple process of trial and error but such a process would contain a strong ad hoc element in the way it models the sequence of trials. Firms do definitely not proceed by purely random trials. When learning is individual[1], new strategies are necessarily based on the past experience: firms combine known decision rules in order to reach better ones. Genetic Algorithms (**GA**), implement such a learning process through evolutionary mechanisms: from a population of actual decision rules, the selection keeps the best ones, the crossover combines these in order to obtain better rules and the mutation introduces some small amount of random experimenting. Moreover, they have the capacity to robustly handle quite complex environments (see Goldberg [1991] for several examples) and, in this sense, can well correspond to the conditions of R&D decisions. One should not conceive GA as a way to represent the exact decision mechanism of firms but as a way to represent the presence of learning and of experience-oriented search processes. The GAs also have the capacity to provide a unified modeling strategy in the vast diversity of mechanisms adopted in models of bounded rationality (one could nearly establish a one-to-one correspondence in the literature between models and modelling strategies). Since they take into account learning, they constitute a good rival of fixed rules that seem to actually constitute the only unifying approach.

This article aims to test the relevance of GA, in opposition with fixed R&D rules, in a simplified version of the evolutionary industry model of Nelson and Winter. The original model is simplified in order to focus on R&D process as the main determinant of industry dynamics. Firms arbitrate at each period between R&D and physical capital when allocating their gross profits to different investment projects. The industry is composed of two types of firms: the **NWFirms** that use a fixed decision rule and **GenFirms** that adjust their R&D investment through a GA. Competition selects, in the long term, the firms that outperforms their competitors: firms can only finance their investments by the profits and they must leave the industry when their physical capital vanishes.

[1] Silverberg & Verspagen [1996] as well as Kwasnicki & Kwasnicka [1992] formalize learning as local individual random experiments combined with industry wide imitation. This learning process has both an individual and a population level components. Our objective is to model learning as a fully individual process.

The relevance of explicit inclusion of learning through GA is tested at two levels. First, at the industry level, these modelling strategies are confronted comparing the performances of four cases composed from 0% to 100% of GenFirms. The results of simulations show that the presence of learning firms leads to higher technological performance as well as to higher welfare. In the second place, from the point of view of individual firms, the properties of heterogenous industries are studied in order to assess the competitive role of learning. In fact, learning allows for the discovery of better R&D strategies but it is also costly: firms must test new strategies that quite often happen to be worst than the actual ones. Only the comparison of performances of both types of firms in the long term can establish the utility of learning strategies. Our results show that GenFirms dominate systematically the industry and this domination flows from their learning. These results are obtained through the simulation methodology already developed in Jonard & Yildizoglu [1998]. This procedure uses the comparison, with non-parametric statistical methods, of the results of batches of simulations instead of the comparison of individual simulations.

Java and Win32 binary versions of this program can be found on the web[2]. This site also contains full documentation in Sun's API format.

The remainder of this article is organized as follows. In section two we present the characteristics of the model. The connection between the genetic algorithm and the learning process also is discussed it this section. Section three is dedicated to the presentation of our methodology and results. Section four concludes.

2 The model

I only emphasize new elements in the model. The intersection with the well known Nelson & Winter [1982] (part V, ch.12) model will first be outlined. A second section will present new dimensions included in this model: capital and R&D investments.

2.1 Characteristics common with Nelson and Winter(1982)

At the beginning of each period, the firm j is characterized by the productivity A_j of it's technology and it's capital stock, K_j. Capital is the only production factor, and the production technology is characterized by fixed input coefficient and constant scale economies. Unit using cost of capital, c, is constant over different production techniques (the unit cost of production is c/A_j). The capital stock depreciates at a rate δ at each period.

Production technics are disembodied. There is no switching cost and the capital can be converted without cost from one technology to another (for a more realistic model with vintage capital, see Silverberg, Dosi & Orsenigo [1988]). This corresponds to a vision of technology based on process innovation. In fact, the innovating firm does not replace its capital stock, but uses it more efficiently. An innovation therefore corresponds to better knowledge of the production process.

2.1.1 Production and profits

Each firm in the industry ($j \in I = \{1...N\}$) produces the same homogenous good with the following production function:

$$Q_j = A_j \cdot K_j. \qquad (1)$$

[2] http://cournot.u-strasbg.fr/yildi/learnind/index.html.

The gross profit rate on capital of the firm is given by:

$$\pi_j = pA - c \tag{2}$$

where p is the market price and is determined by a short term equilibrium on the product market:

$$\begin{cases} Q = \sum_j Q_j \\ p = p(Q) = \dfrac{\mathbf{D}}{Q^{1/\eta}} \end{cases} \tag{3}$$

where Q is the total supply, $p(Q)$ is a constant elasticity inverse market demand function, and η is the Marshalian demand elasticity. Gross profits of the firm are given by

$$\Pi_j = \pi_j \cdot K_j \tag{4}$$

The state of each firm will change from one period to another in consequence of the R&D decisions, which modify its technology and hence its productivity, and the investment behavior, which modifies its capital stock.

2.1.2 R&D and technical progress

The productivities are modified in each period consequently to the technical progress. In each period firms invest RD_{jt} on R&D. This investment allows them to imitate their successful competitors and to innovate. Both imitation and innovation are two-stage stochastic processes.

Innovation

Innovation is a two-stage stochastic process. A first draw determines if the R&D investment of the firm has been successful and resulted in an innovation:

$$P[d_{int} = 1] = a_n \cdot RD_{jt},$$

where a_n is a calibration parameter that projects RD on $[0,1]$. A second draw gives the effective result of the innovation

$$\tilde{A}_{jt} \rightsquigarrow N\left(A_{jt}, \sigma_2^2\right).$$

Hence innovation is a cumulative process and firms with higher productivities have better chance to attain even higher productivities.

Imitation

For the imitation, we have one stochastic draw which determines if the firm's R&D investment has been successful. If it is the case, the firm obtains the best practice in the industry (A_t^*):

$$P[d_{imt} = 1] = a_m \cdot RD_{jt}$$
$$\hat{A}_{jt} = A_{jt} + d_{imt} \cdot \left(A_t^* - A_{jt}\right).$$

New productivity of the firm

Finally, the effective productivity of the firm for the next period is given by the best of these three outcomes:

$$A_{j,t+1} = \max\left\{A_{jt}, \tilde{A}_{jt}, \hat{A}_{jt}\right\} \tag{5}$$

2.2 Capital investment and R&D decisions

Main differences between this model and Nelson & Winter (1982) consist in the investment behaviour: investment in physical capital and investment in R&D. A possibility of exit from the industry is also included in the model. In each period firms invest a fraction of their gross profit on R&D. The rest of this profit is used for the expansion of physical capital.

2.2.1 R&D decisions and genetic algorithms

Firms invest a fraction rd_{jt} of their gross profits on R&D. A minimal investment is necessary to keep *alive* the R&D potential (research equipment and team). We therefore have $rd_{jt} \geq rd_{\min}$.

There are two types of firms: NWFirms and GenFirms. They are distinguished by their R&D investment behaviour.

NWFirms invest in each period a fixed proportion rd_{NW} of their profit in R&D (in addition to the minimal amount of R&D):

$$RD_{NW\,jt} = (rd_{\min} + rd_{NW}) \cdot \Pi_{jt} \tag{6}$$

This rule corresponds to the representation of bounded rationality by "fixed rules". Learning of firms about their environment does not influence their R&D behaviour. This is the common approach retained in many evolutionary industry models. Learning is taken into account in the behaviour of GenFirms.

Each **GenFirm** uses an individual genetic algorithm (GA) in order to adjust the R&D strategy (the fraction $rd_{jt} \geq rd_{\min}$) to the conditions of the industry. Each possible strategy of the firm is coded as a chromosome C_i of length **G**. During its life, the firm carries a population of **C** strategies (number of chromosomes). This population of parallel rules evolves as a consequence of the experience of the firm in the industry.

The experience of the firm can only influence these rules if it provides an evaluation mechanism for different rules. In an industrial context, the only way of evaluating a rule is using it: the value of a rule depends on the dynamics of the industry and hence, on the behaviour of other firms. Moreover, R&D investment does not pay back immediately and each R&D strategy must be used for many periods before proving its efficiency. Consequently, in order to evaluate each rule, the firm uses it for a number of periods (n =*learning period*) and **the average gross profit rate of this time interval gives the fitness of this strategy**. When all strategies of the population are evaluated, a new population is generated through selection–crossover–mutation. We use an elitist GA that conserves the best strategy of the preceding period in the population.

We adopt an indirect coding of R&D strategies: the fraction of profits dedicated to R&D (strategy) is coded as a chromosome C_i of length **G**. The decimal value of the chromosome corresponds to the position of this strategy in the search space $[0\%, 100\%]$. This space contains $\Delta = \sum_{i=0}^{i=G-1} 1 \cdot 2^i$ equally spaced strictly positive strategies, and zero. The R&D rate corresponding to a chromosome C_i is finally computed using the following rule

$$rd_j = (C_i)_{10} \cdot \frac{1}{\Delta} + rd_{\min}. \tag{7}$$

Example: If $G = 4$, there are $(1111)_{10} = 1 \cdot 2^3 + 1 \cdot 2^2 + 1 \cdot 2^1 + 1 \cdot 2^0 = 15 = \Delta$ strictly positive strategies equally spaced between 0% and 100%. If a strategy of the firm is $C_i = 0011$, this chromosome corresponds to the following R&D investment rate:

$$C_i = (0011)_{10} = 3 \Rightarrow rd_j = 3 \cdot \frac{1}{15} + rd_{\min} = 20\% + rd_{\min}.$$

Finally, the R&D investment of the firm is given by

$$RD_{jt} = \left(rd_{\min} + rd_{jt}\right) \cdot \Pi_{jt}$$

Even if the GA does not represent the exact learning mechanism of firms, it is a convenient way of representing the presence of this learning at the individual level. Our representation of the learning process is significantly different from the one considered by Brenner [1998] in his comparison of evolutionary algorithms with learning algorithms. Many limits to which Brenner [1998] draws our attention concerns the use of the evolutionary algorithms at the population level. Quite differently, we use the GA to represent learning of rules at the individual level: each firm carries an individual population evolving of decision rules. Our formulation gives a clear microeconomic foundation to learning in accordance with the modelling of the industry dynamics. The importance of this point is clearly established by Vriend [1998]. This formulation also excludes many ambiguities that appear when one models learning of rules at the industry level: the definition of fitness at the industry level, the connection between selection and performance of individual firms are the most disturbing of them.

Consequently, the selection–crossover–mutation mechanisms respectively correspond, at the firm level, to the elimination of bad rules, to the combination of the selcted rules in order to discover better new rules, and to few random experiments. Elitism assures that memory is taken into account and *good* old rules are not eliminated if better new rules are not found. Also, the chromosome length, **G**, has a signification in terms of the learning process of firms: the higher **G**, the finer the search process of the firm. A firm that uses a higher **G** is more demanding for its learning process: it desires to get closer to the best strategies in the search space. But learning will be costlier for such a firm because it will have to try many rules before getting closer to the best ones. Consequently, the boundaries of the search space do not depend on **G** (firms always explore $[0\%, 100\%]$) but the refinement of the process does. Our coding mechanism hence combines the best of two worlds: the speed of binary coding and the flexibility of *graycode*. A higher number of chromosomes $(\mathbf{C})$ corresponds to a more flexible learning process that will conserve more rules in the rule population of the firm, but this flexibility will also have a cost: the higher the number of rules in the population, the higher the total learning period for each particular population of rules. Consequently, a nice correspondence exists between the characteristics of the GA and the learning processes of firms.

2.2.2 Investment in physical capital

Capital investment results directly from the arbitrage of firms between the R&D investment and capital expansion. Learning firms adapt the sharing of gross profits between R&D and physical capital:

$$K_{jt+1} = (1 - \delta) \cdot K_{jt} + \left(1 - rd_{jt}\right) \cdot \Pi_{jt}.$$

2.2.3 Exit

If the profits of the firm get persistently low, it can loose all possibility of investment and innovation. In this case, current profits do not permit any investments. The capital stock of the firm vanishes because of the depreciation. When the capital stock gets very small, the firm looses all possibility of innovation and growth. It consequently exits the industry when

$$K_{jt} \leq \underline{K}.$$

3 Simulation methodology and results

I use the simulation protocol developed in Jonard & Yildizoglu [1998]. This protocol is explained in a first paragraph. Relative performance of GenFirms is measured through different indicators that have been developed for this article. Simulation results are used to assess the role of learning at two different levels. First, the role of learning on the aggregate performance of the industry is explored. Second, the relative performance of learning firms – i.e. their competitive edge – is evaluated.

3.1 Protocol

Since we aim to derive results independent from a particular sequence of random numbers, a batch of 20 simulations, of 6000 periods each, is run for each configuration of the model. Observations have been saved every 40 periods. The whole possible history of the industry is hence represented by a sample of 3000 observations. The relevant dimensions (e.g. technical progress, concentration) of resulting samples are compared by way of non-parametric testing (the non-parametric Wilcoxon–Mann–Whitney test, see for instance ch.18 in Watson, Billingsley, Croft & Huntsberger [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.

3.2 Indicators

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

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

Only the concentration of capital needs some explanation:

$$K = \frac{\sum_j K_j^2}{\left(\sum_j K_j\right)^2}, \tag{8}$$

where K_j is the capital stock of firm j. This indicator gives an equivalent number of firms as if each of them had the same part of capital stock. We have $1 \leq K \leq N$ where N is the number of active firms in the industry. The higher is this indicator, the more evenly balanced is the distribution of capital stock between firms. This is an application of the Herfindall index to the capital stocks and summarizes the inequalities in the distribution of the capital stock.

Some simple new indicators are necessary in order to compare relative performance of GenFirms:

- the share of capital owned by GenFirms;
- the share of GenFirms in cumulated profits;
- market share of GenFirms;
- R&D investment share of GenFirms.

Since the shares of NWFirms are complementary and give a total of 100%, only the shares of GenFirms are used for comparison.

3.3 Comparison of industry performance

Three different points of view can be adopted for the evaluation of the impact of learning on industrial efficiency: technological performance, firms' profit and consumers' welfare. We do not have a direct indicator for consumers' welfare, but the market price is of course inversely related to consumers' surplus. The effect on firms' surplus can be appreciated by comparing the distribution of average cumulative profits in each industry configuration. Technological performance is evaluated through average and maximal productivity. The latter shows how far a particular industry can go in the technology space and the former resumes general technological level of industry.

The presence of learning firms should normally increase technological efficiency because these firms are able to exploit the increasing relationship that exists between R&D investment and innovation. But there is a specific cost for learning: in order to learn, firms must spend time to try different strategies, including the inferior ones. Learning can consequently be a source of delay in the discovery of better technologies. The overall effect can only assessed through the comparison of different industries.

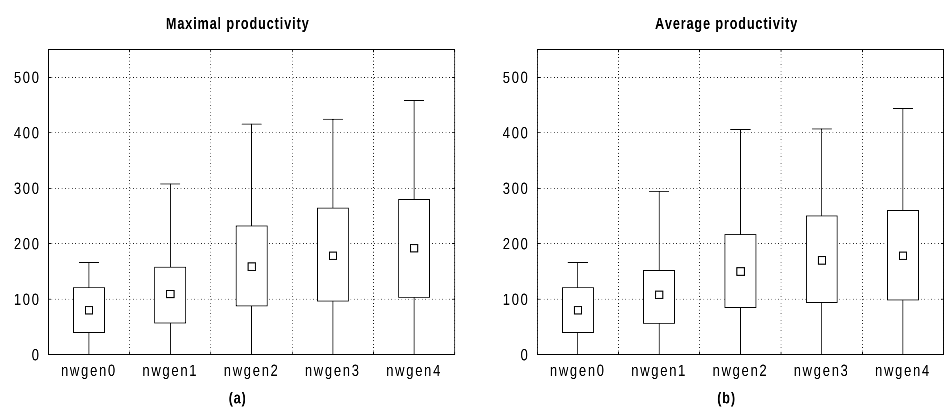

Figure 1: Learning and technology dynamics

We compare four industry configurations:

nwgen0: 100% of NWFirms;

nwgen1: 75% of NWFirms, 25% of GenFirms;

nwgen2: 50% of each type;

nwgen3: 75% of GenFirms;

nwgen4: 100% of GenFirms.

All configurations have a total population of 40 firms and all GenFirms are the simplest kind, they have $C = 8$ chromosomes of $G = 7$ genes. NWFirms invest $(rd_{\min} + 7\%)$ in R&D. $rd_{\min} = 3\%$. Other parameters are common to all simulations and they are given in appendix.

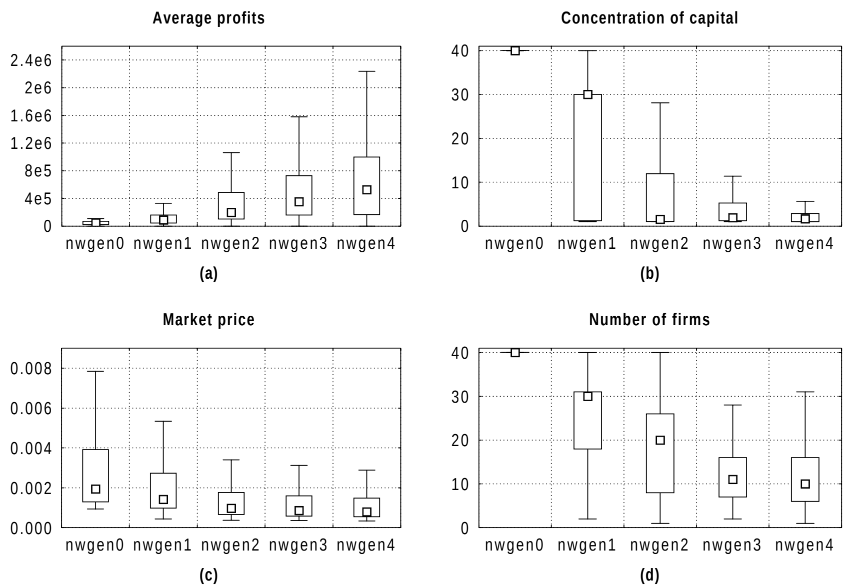

Figure 2: Learning and market structure

The results on technological efficiency are represented in Figure 1. Graphic (a) clearly shows that learning firms contribute very significantly to the technological advancement of the industry. Their impact is important even if they form a minority (even 25%). Moreover, the difference with the distribution of average productivity (Graphic (b)) is very small and consequently the diffusion is very quick in these industries. Higher technological efficiency is due to the presence of learning. This is summarized in the following proposition.

Proposition 1 *Presence of learning firms is a source of technological efficiency at the industry level.*

One could, wrongly, be tempted to explain this positive effect by the low *rd* ratio of NWFirms, but we have very similar distributions even when $rd_{NW} = 27\%$ (see Figure 4).

Quite interestingly, this efficiency is even costless for society. The presence of learning firms increases the concentration of capital (equivalent number of firms decreases in Figure 2–(b)) but this higher concentration does not increase the market price (Figure 2–(c)). Hence the effect on consumers' welfare is not negative (Figure 2–(c)). Moreover, higher investments in R&D do not even penalize the gross profits of the firms (Figure 2–(a)): learning is even a source of supplementary profits for the industry and the global effect of learning on society is clearly positive.

Proposition 2 *Presence of learning firms implies*

1. *higher concentration;*

2. *higher gross profits;*

3. *lower market price;*

4. *higher social welfare.*

These consequences clearly result from the evolution of the arbitrage of GenFirms between R&D investment and capital investment. This proposition also implies that if we neglect learning, we can overestimate the welfare loss generally associated to greater concentration: even **nwgen1** (25% of learning firms) clearly improves the social welfare in comparison with **nwgen0** (100% of NWFirms). Learning firms deliberately modify both components of their production: cost and capital stock. The presence of learning firms is hence a source of dynamic social efficiency at the industry level and the efficiency at the technological level is the real source of this positive effect on social welfare.

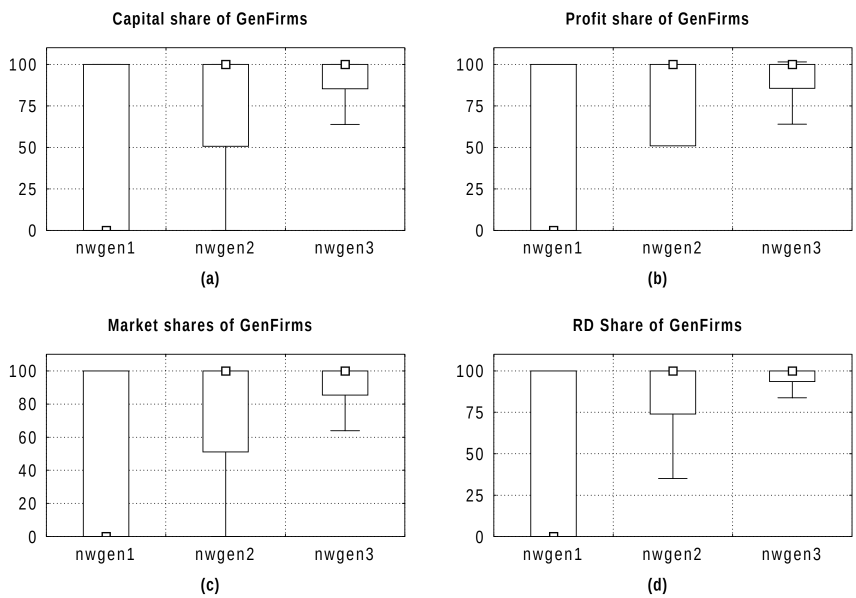

Figure 3: Relative performance of GenFirms

3.4 Competing R&D strategies

Relative performance of GenFirms can be measured by their share in aggregate magnitudes of the industry. It is indeed important to show that learning firms are effectively benefiting from this learning. We use hybrid industries (**nwgen1, nwgen2** and **nwgen3**) for this comparison.

The Figure 3 clearly shows that when GenFirms compose more than 50% of the initial population, they dominate the whole history of the industry (Figures 3–(a)-(c)). This domination comes from a higher investment on R&D than the NWFirms (Figure 3–(d)).

Proposition 3 *When they do not form a* too *small minority, learning firms dominate the market and gain shares comparatively to their initial positions.*

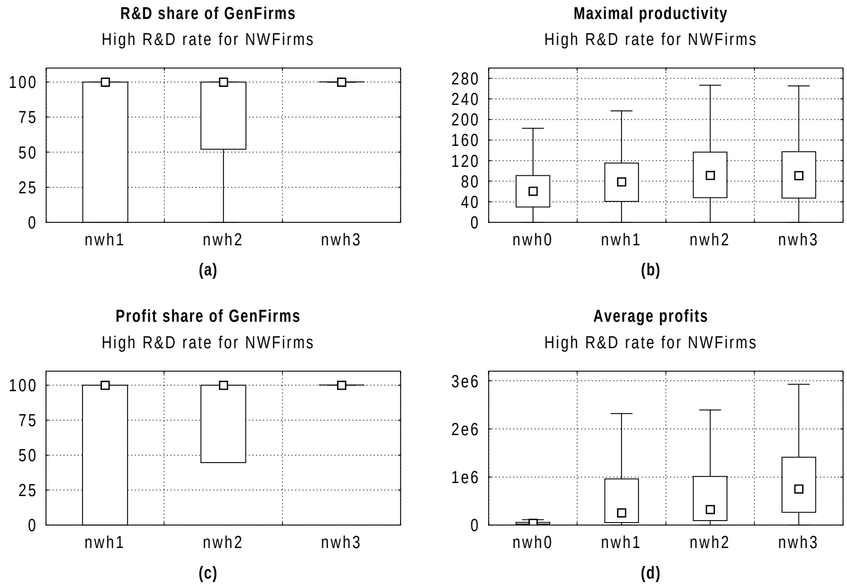

Figure 4: $rd_{NW} = 27\% + 3\%$

The costly and random nature of learning plays against learning firms when they are too few in the industry and their *mistakes* eventually push them out of the industry. When they are numerous, some of them succeed and they end up dominating the industry. One could think that the low R&D rate of NWFirms is responsible of this result. The following figure shows that even with higher fixed R&D rate, NWFirms cannot dominate the industry.

Proposition 4 *Even if the R&D rate of NWFirms is high, GenFirms dominate the R&D activity of the industry.*

The Figures 4–(a) and (c) show that even when facing NWFirms with higher R&D rate, GenFirms dominate the industry. In fact, higher R&D rate imposes a stringent constraint on capital investment of NWfirms while GenFirms are continuously arbitrating between these two investments. Their relative performance is even higher in this case. Figures 4–(b) and (d) again indicate the positive impact of learning firms on the performances of industry.

In order to check the reality of learning, we need to abandon our methodology and consider an individual simulation (the last of the 20 simulations). The Figure 5 gives central indicators of the distribution of R&D rate of the GenFirms in **nwgen2**. We represent in this figure $(\mu - \sigma, \mu, \mu + \sigma)$ where μ is the average and σ is the standard deviation. This Figure clearly shows that GenFirms are not simply randomizing and the dispersion is decreasing in time.

Proposition 5 *GenFirms learn.*

4 Conclusion

This article is a first attempt to explicitly compare different behaviour rules for R&D investment. Ballot & Taymaz [1999] have already done such a comparison but the complexity of the underlying

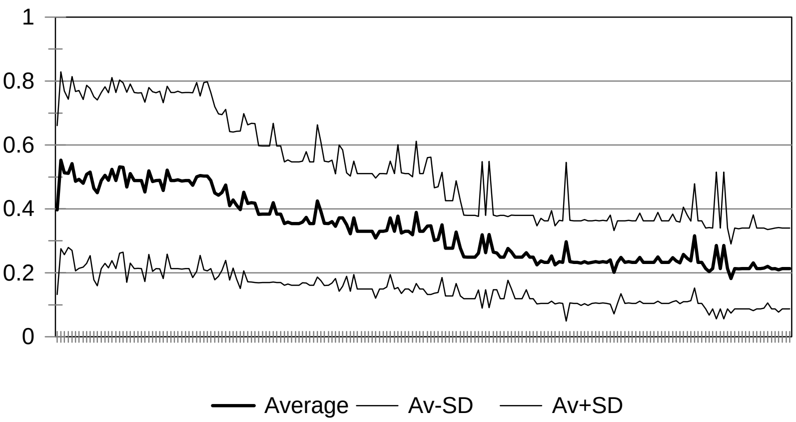

Figure 5: Convergence of rd_{jt} in **nwgen4**

model (MOSES model of Swedish industry) considerably conceals the exact role of different decision rules in industry and firms performance. I deliberately adopt a very simple model in order to completely focus on the effects of R&D rules.

Two general results dominate the simulations. In the first place, results at the industry level clearly show that we should not ignore learning in models of industry. Otherwise, this can result in a severe underestimation of the performance of industries at the technological level and, at the level of social welfare: industries with learning firms exhibit higher technological and social efficiency. The imperfect competition generally associated to the innovation process is not necessarily the cause of a significant loss of welfare, even in the short term. In the second place, learning gives a competitive edge to firms benefitting from it: learning firms dominate the industry. Both results are directly engendered by the continuous arbitrage of learning firms between R&D and capital investments.

On a methodological level, one of the shortcomings of Genetic Algorithms in industrial context with endogenous payoff structure is the necessity of effectively using each rule in order to discover its fitness. Learning is consequently slow (a different but similar problem also applies to classifier systems): firm's learning is directly on the strategy space. A more ambitious assumption about learning would consider firms that aim to discover as much of information as possible on the payoff structure; to have *expectations* on the relationship between R&D and profit. Such a learning would be based on inductive reasoning (Holland, Holyoak & Thagard [1989]). Oltra & Yildizoglu [1998] propose to model these expectations using an artificial neural network (**ANN**). A more complete learning model should then proceed in cascade: a GA searching the strategy space and an ANN providing expected fitness values for strategies. Learning would in this case include a better understanding of the environment of the firm (through the adjustment of the ANN) and the discovery of better strategies (through the workings of the GA) given this understanding.

References

Ballot, G. & 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**(2).

Brenner, T. [1998], 'Can evolutionary algorithms describe learning processes?', *Journal of Evolutionary Economics* (8), 271–283.

Goldberg, D. E. [1991], *Genetic Algorithms*, Addison-Wesley, Reading: MA.

Holland, J. H., Holyoak, K. J. & Thagard, P. R. [1989], *Induction. Processes of Inference, Learning, and Discovery*, MIT Press, Cambridge:MA.

Jonard, N. & Yildizoglu, M. [1998], 'Technological diversity in an evolutionary industry model with localized learning and network externalities', *Structural Change and Economic Dynamics* **9(1)**, 35–55.

Knight, F. H. [1921], *Risk, Uncertainty and Profits*, number Reprint, Chicago University Press, Chicago.

Kwasnicki, W. & Kwasnicka, H. [1992], 'Market, innovation, competition. an evolutionary model of industrial dynamics', *Journal of Economic Behavior and Organization* **19**, 343–368.

Nelson, R. R. & Winter, S. [1982], *An Evolutionary Theory of Economic Change*, The Belknap Press of Harvard University, London.

Oltra, V. & Yildizoglu, M. [1998], 'Expectations and adaptive behaviour: the missing trade-off in models of innovation', *mimeo,BETA. Universite Louis Pasteur, Strasbourg* .

Silverberg, G., Dosi, G. & Orsenigo, L. [1988], 'Innovation, diversity and diffusion: a self-organization model', *Economic Journal* **98**, 1032–1054.

Silverberg, J. & 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 behavoristic model, in Bowman, M. J. (ed), Expectations, Uncertainty and Business Behavior, Social Science Council, New York, pp. 49–58.

Simon, H. A. [1976], From substantial to procedural rationality, in Latsis, S. J. (ed), Method and Appraisal in Economics, Cambridge University Press, Cambridge, pp. 129–148.

Vriend, N. [1998], 'An illustration of the essential difference between individual and social learning, and its consequences for computational analyses', *Journal of Economic Dynamics and Control* **(forthcoming)**.

Watson, C. J., Billingsley, D. J., Croft, D. J. & Huntsberger, D. V. [1993], *Statistics for Management and Economics*, fifth edition, Allyn and Bacon, Boston.

Appendix: parameter values

d_{in} is fixed in order to have a initial innovation probability of 5%. $d_{im} = d_{in}/10$.

Parameter	Value
Number of NWFirms: N_{NW}	$variable$
Number of GenFirms: N_{Gen}	$variable$
Output frequency	40
Number of simulations	20
Number of periods: T	3000
Using cost of capital: c	0.1
Initial productivity: A_0	0.16
Initial capital: K_0	50
Demand elasticity: η	0.5
Autonomous demand: $\mathbf{D}$	100
Depreciation rate: δ	5%
Threshold capital: $\underline{K}$	10^{-5}
R&D rate of NWFirms: rd_{NW}	7%
Minimal R&D rate: $rd_{\min}$	3%
Dispersion of Innovations: σ	0.05
Number of chromosomes: C	8
Number of genes: G	7
$Xover\ points$	1
Learning rate: n	5
$P[Xover]$	0.7
$P[Mutate]$	0.03