



HAL
open science

Cournot Competition and Endogenous Firm Size

Jason Barr, Francesco Saraceno

► **To cite this version:**

Jason Barr, Francesco Saraceno. Cournot Competition and Endogenous Firm Size. 2005. hal-01052859

HAL Id: hal-01052859

<https://sciencespo.hal.science/hal-01052859>

Preprint submitted on 28 Jul 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Cournot Competition and Endogenous Firm Size*

Jason Barr

Rutgers University, Newark

jmbarr@rutgers.edu

Francesco Saraceno

Observatoire Français des Conjonctures Économiques

francesco.saraceno@sciences-po.fr

January 2005

Rutgers University Newark Working Paper #2005-001

Abstract

We study the dynamics of firm size in a repeated Cournot game with unknown demand function. We model the firm as a type of artificial neural network. Each period it must learn to map environmental signals to both demand parameters and its rival's output choice. But this learning game is in the background, as we focus on the endogenous adjustment of network size. We investigate the long-run behavior of firm/network size as a function of profits, rival's size, and the type of adjustment rules used.

Keywords: Firm size, adjustment dynamics, artificial neural networks, Cournot games

JEL Classification: C63, D21, D83, L13

*An earlier version of this paper was presented at the 10th International Conference on Computing in Economics and Finance, July 8-10, 2004, University of Amsterdam. We thank the session participants for their comments.

1 Introduction

In this paper we explore firm size dynamics, with the firm modeled as a type of artificial neural network (ANN). Two firms/networks compete at two different levels. The first level, which has been explored in detail in other work (Barr and Saraceno (BS), 2004; 2005), looks at Cournot competition between two neural networks. In this paper, this level of competition is essentially in the background, while the main form of strategic interaction is in regards to firm size dynamics. The firm, while playing the repeated Cournot game, has to make long run decisions about its size, which affects not only its own profits, but those of its rival as well.

Our previous research showed that firm size, which was left exogenous, is an important determinant of performance in an uncertain environment. Here we reverse the perspective, taking as given both the learning process and the dependence of firm profit on size and environmental complexity, and endogenize firm size in order to investigate whether simple adjustment rules succeed in yielding the optimal size (defined as the result of a best response dynamics). The computational requirements needed to discover the optimal network size may be quite expensive for the firm; and thus we explore less costly methods for adjusting firm size.

In particular, we explore two types of adjustment rules. The first rule ("the isolationist") has the firm adjusting its size simply based on its last period profit growth. The second rule ("the imitationist") has the firm adjusting its size if its rival has larger profits. Lastly we investigate firm dynamics when the firms use a combined rule.

We explore how the two firms interact to affect their respective dynamics; we also study how the adjustment rate parameters affect dynamics and long run firm size. In addition, we investigate the conditions under which the adjustment rules will produce the Cournot-Nash outcome.

To our knowledge, no other paper has developed the issue of long run firm growth in an adaptive setting. Here we highlight some of the important findings of our work:

- In the isolationist case, firm dynamics is a function of both environmental complexity, initial size, rival's initial size and the adjustment rate parameter. The value of a firm's long run size is a non-linear function of these four variables. **Ceteris paribus**, we find, for example, that firms using a simple rule gives decreasing firm size versus increasing

complexity. In fact, a complex environment causes profit growth, and consequently, firm growth, to be lower.

- In the imitationist case, firm dynamics is also a function of initial firm size, initial rival's size, complexity and the adjustment rate parameter. In such a case, increasing complexity will yield lower steady state sizes, though they are not necessarily the same for the two firms.
- Via regression analysis we measure the relative effects of the various initial conditions and parameters on long run dynamics. We find that own initial size is positively related to long run size, rival's initial size has a negative effect for small initial size, but positive effect for larger initial sizes. The larger the adjustment parameters, the larger is the long run size.
- Finally, we show that the simple dynamics that we consider very rarely converge to the 'best response' outcome. In particular, the firm's adjustment parameter plays an important role in guaranteeing such a convergence.

Our work relates to a few different areas. Our approach to using a neural network fits within the agent-based literature on information processing (IP) organizations (Chang and Harrington, forthcoming). In this vein, organizations are modeled as a collection or network of agents that are responsible for processing incoming data. IP networks and organizations arise because in modern economies no one agent can process all the data, as well as make decisions about it. The growth of the modern corporation has created the need for workers who are managers and information processors (Chandler, 1977; Radner, 1993).

Typical models are concerned with the relationship between the structure of the network and the corresponding performance or cost (DeCanio and Watkins, 1998; Radner, 1993; Van Zandt, 1998). In this paper, the network is responsible for mapping incoming signals about the economic environment to both demand and a rival's output decision. Unlike other information processing models, we explicitly include strategic interaction: one firm's ability to learn the environment affects the other firm's pay-offs. Thus a firm must locate an optimal network size not only to maximize performance from learning the environment but also to respond to its rival's actions. In our

case, the firm is able to learn over time as it repeatedly gains experience in observing and making decisions about environmental signals.

A second area of literature that relates to our work is that of the evolutionary and firm decision making models of Nelson and Winter (1982), Simon (1982) and Cyret and March (1963). In this area, the firm is also boundedly rational, but the focus is not on information processing *per se*. Rather, the firm is engaged in a myriad of activities from production, sales and marketing, R&D, business strategy, etc. As the firm engages in its business activities it gains a set of capabilities that cannot be easily replicated by other firms. The patterns of behavior that it collectively masters are known as its 'routines' (Nelson and Winter, 1982). Routines are often comprised of rules-of-thumb behavior: continue to do something if it is working, change if not.

In this vein, firms in our paper employ simple adjustment rules when choosing a firm size. In a world where there is an abundance of information to process, and when discovering optimal solutions is often computationally expensive firms will seek relatively easier rules of behavior, ones that produce satisfactory responses at relatively low cost (Simon, 1982).

The rest of the paper is organized as follows. Section 2 discusses the motivation for using a neural network as a model of the firm. Then, in section 3, we give a brief discussion of the set up of the model. A more detailed treatment is given in the Appendix. Next, section 4 gives the benchmark cases of network size equilibria. Sections 5 to 7 discuss the heart of the paper—the firm size adjustment algorithms and the results of the algorithms. Finally, section 8 presents a discussion on the implications of the model and also gives some concluding remarks.

2 Neural Networks as a Model of the Firm

In previous work (BS, 2002; 2005) we argued that information processing is a crucial feature of modern corporations, and that efficiency in performing this task may be crucial for success or failure. We further argued that when focussing on this aspect of firm behavior, computational learning theory may give useful insights and modelling techniques. In this perspective, it is useful to view the firm as a **learning algorithm**, consisting of agents that follow a series of rules and procedures organized in both a parallel and serial man-

ner. Firms learn and improve their performance by repeating their actions and recognizing patterns (i.e., learning by doing). As the firm processes information, it learns its particular environment and becomes proficient at recognizing new and related information.

Among the many possible learning machines, we focussed on Artificial Neural Networks as models of the firm, because of the intuitive mapping between their parallel processing structure and firm organization. Neural networks, like other learning machines, can generalize from experience to unseen problems, i.e., they recognize patterns. Firms (and in general economic agents) do the same: the know-how acquired over time is used in tackling new, related problems.

What is specific about ANNs, as learning machines, is the parallel and decentralized processing. ANNs are composed of multiple units processing relatively simple tasks in parallel. The combined result of this multiplicity is the ability to process very complex tasks. In the same way firms are often composed of different units working autonomously on very specific tasks, and are coordinated by a management that merges the results of these simple operations in order to design complex strategies.

Furthermore, the firm, like learning algorithms, faces a trade-off linked to the complexity of its organization. Small firms are likely to attain a rather imprecise understanding of the environment they face; but on the other hand they act pretty quickly and are able to design decent strategies with small amounts of experience. Larger and more complex firms, on the other hand, produce more sophisticated analyses, but they need time and experience to implement their strategies. Thus, the optimal firm structure may only be determined in relation with the environment, and it is likely to change with it. Unlike computer science, however, in economics the search for an optimal structure occurs given a competitive landscape, which imposes time and money constraints on the firm.

In our previous work we showed, by means of simulations, that the trade-off between speed and accuracy generates a hump-shaped profit curve in firm size (BS, 2002). We also showed that as complexity of the environment increases the firm size that maximizes profit also increases. These results reappeared when we applied the model to Cournot competition. Here, we leave the Cournot competition and the learning process in the background, and investigate how network size changes endogenously.

3 Setup of the Model

The setup of the neural network model is taken from BS (2004) and BS (2005). Two firms competing in quantities face a linear demand function whose intercept is unobserved. They observe a set of environmental variables that are related to demand, and have to learn the mapping between the two. Furthermore, firms have to learn their rival's output choice. Appendix B gives the details of the model. BS (2005) shows that in general firms are capable of learning how to map environmental factors to demand, which allows them to converge to the Nash equilibrium. We further showed that the main determinants of firm profitability are, on one hand, firm sizes (i.e., the number of processing units of the two firms, m_1 and m_2); and on the other environmental complexity, which we modeled as being the number of inputs to the network (n), i.e., the number of environmental factors affecting demand.

These facts may be captured by a polynomial in the three variables:

$$\pi_i = f(m_1, m_2, n) \quad i = 1, 2. \quad (1)$$

To obtain a specific numerical form for equation 1, we simulated the Cournot learning process with different randomly drawn firm sizes ($m_1, m_2 \in [2, 20]$) and complexity ($n \in [5, 50]$), recording each time the profit of the two firms. With this data set we ran a regression, which is reported in table 4 in the appendix (notice that the setup is symmetric, so that either firm could be used).

The specific polynomial relating profits to size and complexity, which will serve as the basis for our firm size dynamics, is given in the following equation:¹

$$\begin{aligned} \pi_1 = & 271 + 5.93m_1 - 0.38m_1^2 + 0.007m_1^3 + 0.49m_2 \\ & - 0.3m_1m_2 - 2.2n + 0.0033n^2 + 0.007m_1m_2n - 0.016m_2n. \end{aligned} \quad (2)$$

Note that we do not directly consider the costs of the firm's network since under standard assumptions they do not affect the qualitative results. Figure 1, shows the hump shape of profit with respect to own size, that we discussed

¹In previous work other variables affected profit, here we hold them constant. Furthermore, the regression results are reported for $10,000 \cdot \pi_1$ to make the equation easier to read.

in previous work. Three curves are reported, corresponding to small, medium and large opponent's size (complexity is fixed at $n = 10$).

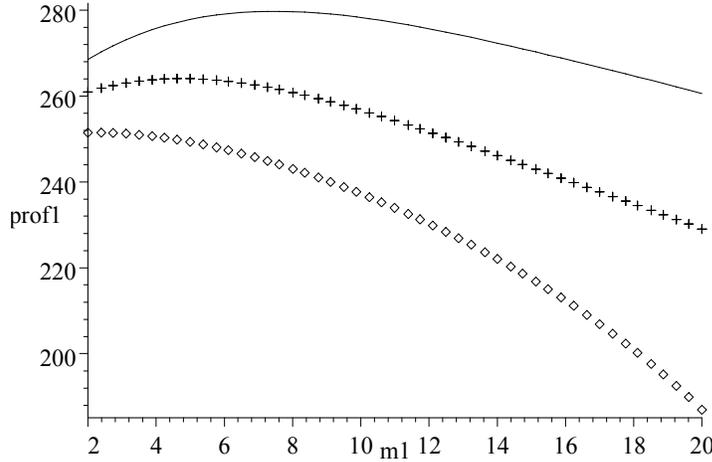


Figure 1: A firm's profit function vs. own size ($m_2 = 2$, solid line; $m_2 = 10$, crosses; $m_2 = 20$, diamonds). $n = 10$.

4 The Best Response Function and Network Size Equilibrium

In this section, as a benchmark, we discuss the firm's best response function and equilibria that arise from the game. We can derive the best response function in size by setting the derivative of profit with respect to size equal to zero, i.e.,

$$\frac{\partial \pi_i}{\partial m_i} = f'(m_i, m_{-i}, n) = 0 \quad i = 1, 2. \quad (3)$$

Given the functional form for profit of equation (2), this yields the following solution for firm 1 (as usual either firm can be used, as the problem

is symmetric)²:

$$m_1^{br}(m_2, n) = 16.9 \pm 2.26\sqrt{2.6m_2 - 0.058nm_2 + 3.9}$$

The 'best response' function is polynomial in m_i and as a result, there is generally more than one solution, and often some of the solutions are complex numbers. Nevertheless, for values of m_2 and n in the admissible range ($m_i \in [2, 20]$, $n \in [5, 50]$), the solution is unique and decreasing. The Network Size Equilibrium (NSE) is given by the intersection of the best responses for the two firms. Figure 2 shows the best response mappings for equation (2), and the corresponding Nash equilibria in size; notice that these equilibria are stable. In fact, in spite of the complexity of the profit function, the best response is quasi-linear. Notice further the relationship between the best responses and environmental complexity: increasing complexity shifts the best response functions, and consequently the Nash equilibrium upwards. In

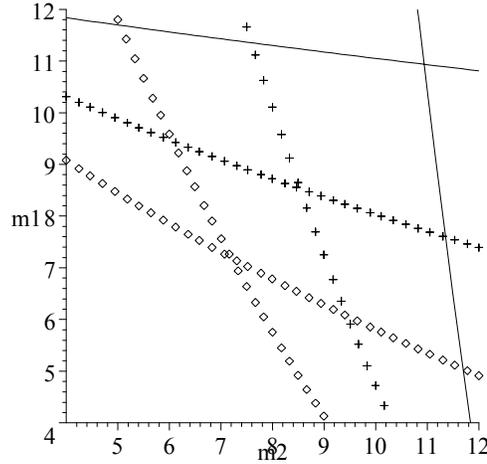


Figure 2: Best response functions ($n = 10$, diamonds; $n = 25$, crosses; $n = 40$, solid lines) for firms 1 and 2. The Network Size Equilibria are given by the intersection of the lines.

general, we can express optimal firm size, m_i^* , as a function of environmental

²For simplicity we ignore the integer issue in regards to firm size and assume that firm size can take on any real value greater than zero.

complexity. If we take the two best responses, and we impose $m_1^* = m_2^* = m^*$ by symmetry, we obtain the following:

$$m^* = 23.5 - 0.15n - 4.5 \sqrt{\frac{p}{(14.1 - 0.34n + 0.001n^2)}} \quad (4)$$

The plot of this expression may be seen in figure 3.³

Finally, we can ask whether profits are increasing or decreasing at the equilibrium firm size. To answer this question, we substitute the optimal value given by equation (4) into the profit equation (2), to obtain a decreasing relationship; the plot is also reported in figure 3.

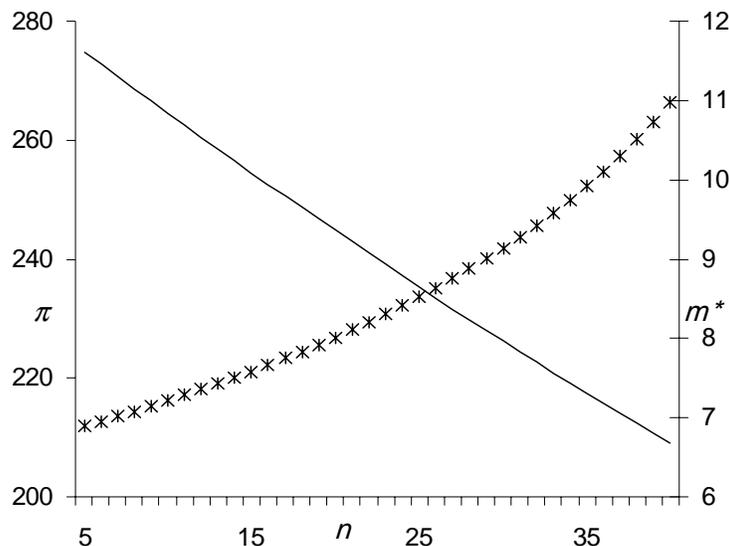


Figure 3: Profit (solid line, left axis) at equilibrium is decreasing in environmental complexity. On the other hand, equilibrium size (crosses, right axis) is increasing.

To conclude, we have shown that best response dynamics yield a unique and stable equilibrium. Furthermore, we were able to show that the firm size in equilibrium is increasing in complexity, while profit is decreasing. In the next section we turn to simpler firm size adjustment rules, that require lower computational capacity for the firm.

³As was the case before, we actually have two solutions for firm size. Nevertheless, one root gives values that are outside the relevant range for m^* , and can be discarded.

5 Adaptive Adjustment Dynamics

As discussed above, firms often face a large amount of information to process. In standard economic theory, firms are assumed to understand many details about how the world functions and how various variables interact. They are assumed to know their cost functions, profit functions, and the effect of a rival's decisions on profits. But in a complex world, with boundedly rational agents, the cost to discover such knowledge is relatively high. As a result, firms learn as they engage in their activities and use this knowledge to help guide them.

For this reason, in this section, we explore relatively simple dynamics for firm size: dynamics that assume on the part of firms simply being able to observe the effect that changing the number of agents (nodes) has on profits. The best response function in section 4 is quite complex, and assumed that the firm knows it; more specifically, it is assumed to know the expected maximal profit obtainable for the entire range of a rival's choice of network size.⁴ (Note also that the best response functions graphed above are numerical approximations.)

In addition in a world in which production and market conditions constantly change, past information may quickly become irrelevant. Meaning that even if a firm has perfect knowledge of its best response function at a certain point in time, that function may quickly become outdated. That means that even when the firm is in possession of the computational capabilities necessary to compute the best response, a firm may not find it efficient to actually do so.

Using the profit function generated in section 3, we explore adjustment dynamics for firms using rule-of-thumb type adjustment rules with the following general adjustment dynamics:

$$m_{i,t} = m_{i,t-1} + \beta (\pi_{i,t-1} - \pi_{i,t-2}) + \alpha I_i [(m_{-i,t-1} - m_{i,t-1})(\pi_{-i,t-1} - \pi_{i,t-1})]. \quad (5)$$

First, β represents the sensitivity of firm size to own profit growth. In other words, if a firm has positive profit growth it will increase its firm size by $\beta (\pi_{i,t-1} - \pi_{i,t-2})$, units; if profit growth is negative, it will decrease by $\beta (\pi_{i,t-1} - \pi_{i,t-2})$. Again, we assume that firm size can take on a real-valued

⁴A common assumption in game theory is that firms know the equilibrium and that they know that their rival knows it, and they know that their rival knows they know it, ad infinitum (Fudenberg and Tirole, 1996).

positive number.

Next, the parameter α captures the "imitation" factor behind size adjustment; I_i is an indicator function taking the value of **1** if the opponent's profit is larger than the firm's, and a value of **0** otherwise:

$$I_i = \begin{matrix} \frac{1}{2} \\ \mathbf{1} \Leftrightarrow (\pi_{-i,t-1} - \pi_{i,t-1}) > \mathbf{0} \\ \mathbf{0} \Leftrightarrow (\pi_{-i,t-1} - \pi_{i,t-1}) \leq \mathbf{0} \end{matrix}$$

In case $I_i = \mathbf{1}$, then size will be adjusted in the direction of the opponents'. Thus, the firm will adjust towards the opponent's size, whenever it observes a better performance of the latter (note that we do not have any discounting).

To sum up, our adjustment rule only uses basic routines: first, the firm expands if it sees its profit increasing; second it adapts towards the opponent's size whenever it sees that the latter is doing better. These are the most basic and commonsensical routines a firm would employ, and require very little observation and computation on the part of the firm. In short, we can think of the Cournot game as happening on a short term basis, while the adjustment dynamics occurs over longer periods of time.

In the next section we investigate the following questions: what kinds of firm dynamics can we expect to see given the simple adjustment rules? And under what parameter choices using equation (5) will firms reach the equilibrium level presented in figure 3 and, what kinds of behavior can we expect for firm size as a function of the parameter space?

6 Results

6.1 Scenario 1: The Isolationist Firm

Suppose that $\alpha = \mathbf{0}$. Then each firm will only look at its own past performance when deciding whether to add or to remove nodes:

$$m_{i,t} = m_{i,t-1} + \beta [\pi_{i,t-1}(m_{i,t-1}, m_{-i,t-1}, n) - \pi_{i,t-2}(m_{i,t-2}, m_{-i,t-2}, n)]$$

Of course, this does not mean that the firm's own dynamics is independent of the other, as in fact $\pi_{i,t-1}$ and $\pi_{i,t-2}$ depend on both sizes. Figure 4 shows the adjustment dynamics of two isolationist firms for different complexity levels. The adjustment rate β is kept fixed (at $\beta = \mathbf{0.05}$).

The dynamics are relatively simple and show a few things. The first is that the long-run level, in general, is reached fairly quickly. The second is

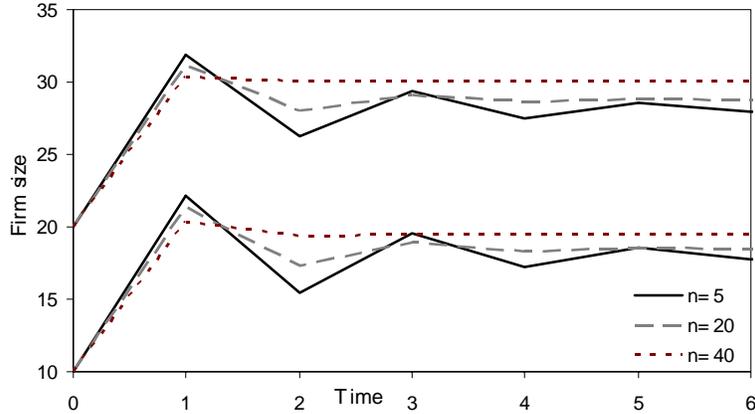


Figure 4: Firm dynamics when the firms start at different sizes ($m_1(0) = 10$; $m_2(0) = 20$), for three different levels of complexity.

that this level seems to depend on initial own size. For example, a firm starting at say 10 nodes will converge at a value around 17-18 nodes; both the opponent's initial size, and the complexity level seem to have a very limited effect, if any. These qualitative features do not change for different initial conditions (other figures available upon request). This result is hardly surprising regarding the opponent's size if we consider that the two firms have a simple adjustment rule, so that interaction only takes place indirectly, through the profit function.

Thus, at first sight the picture does not show strong differences in long run dynamics. However a deeper look will show that initial size, initial rival's size complexity and β all have statistically significant effects on long run size. In section 6.3 below we show regression results for the dynamics of equation (5), where it appears that n and β are important determinants of long run size. Thus, we conclude that in figure 4 the effects of these parameters are hidden by the predominant effect of initial size.

Finally, figure 5 shows the adjustment dynamics for three different values of β . When $\beta = 0.025$, firm size adjusts to approximately 15, when $\beta = 0.075$, long run firm size is larger, finally when $\beta = 0.125$, we see that long run firm size is the largest (about 25) and that there are oscillatory dynamics, with decreasing amplitude, but with apparently regular frequency. Thus we can see that a larger β leads to larger long run firm size.

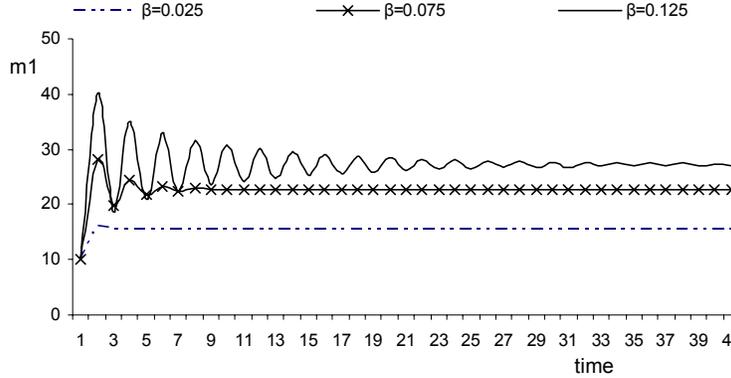


Figure 5: $m_1(0) = m_2(0) = 10$, $n = 20$, three different β .

6.2 Scenario 2: The Imitationist

In this case the contrary of the isolationist firm holds: firms do not care about their own situation, but rather about the comparison with the opponent: $\beta = 0$. Thus,

$$m_{i,t} = m_{i,t-1} + \alpha I_i [(m_{-i,t-1} - m_{i,t-1})(\pi_{-i,t-1} - \pi_{i,t-1})].$$

In this case, the firm will not change its size if it has a larger profit, and it will adjust towards the opponent if it has a smaller profit. Thus, we can say a number of things before feeding actual parameter values. First, at each period, only one firm moves. Second, at the final equilibrium, the two profits must be equal, which happens when firm sizes are equal, but not necessarily only in this situation. For $n = 5$, suppose one firm begins small and the other large ($m_1 = 4$ and $m_2 = 15$). The resulting dynamics depend on α , as shown in figure 6.

For low levels of α , in fact, the drive to imitation is not important enough, and the two firms do not converge to the same size. For intermediate values (around $\alpha = 0.125$), instead, convergence takes place to an intermediate size. When α is too large, on the other hand, the initial movement (in this case of firm 1) is excessive, and may overshoot.

Complexity has a role as well, as shown in figure 7 (where $n = 40$). In fact, as greater complexity yields larger firm size, we see that the adjustment is faster (for each given α), and that for large enough values of α , the system explodes.

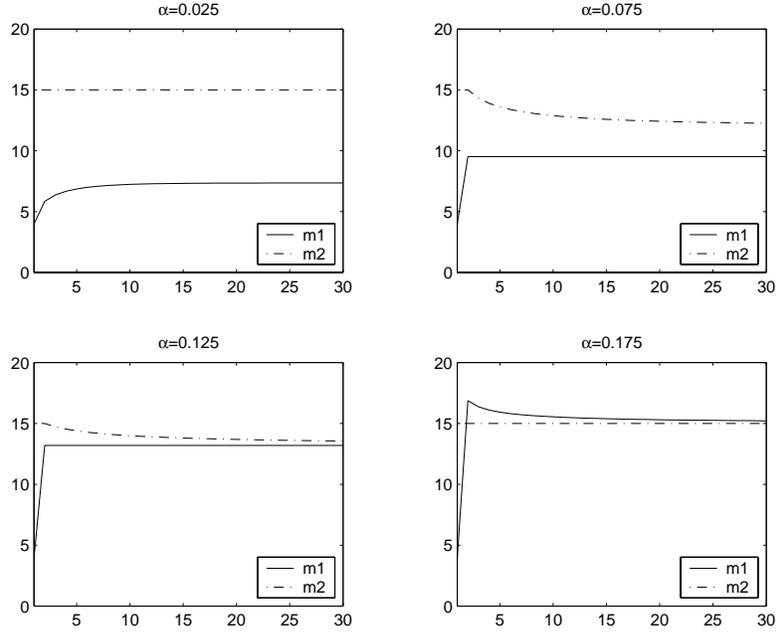


Figure 6: Firm dynamics for different values of α . $n = 5$.

Both in the imitationist and in the isolationist cases, the dynamics show a very strong dependence on initial conditions. This feature of the time series calls for a systematic analysis of the parameter space and we present below regression results for the combined scenario.

6.3 Scenario 3: Combined Dynamics Regression Analysis

In this section, we investigate via regression analysis the effects of initial firm size, adjustment parameters and complexity on long run firm size. Here we combine the two types of dynamics, exploring the outcomes based on equation 5. To do this we generate **5,000** random combinations of initial firm sizes from $m_i(0) \in \{2, 20\}$, $\alpha \in [0.025, 0.075]$, $\beta \in [0.025, 0.075]$ and $n \in \{5, 10, \dots, 40\}$ and then we look at the long run steady state firm size for firm 1, which ranges from **2.6** to **39.6** nodes. The regression results are presented in table 1. The regression equation is non-linear in the variables and includes several interaction terms. Also a dummy variable, $\delta[m_1(0) > m_2(0)] = 1$ if

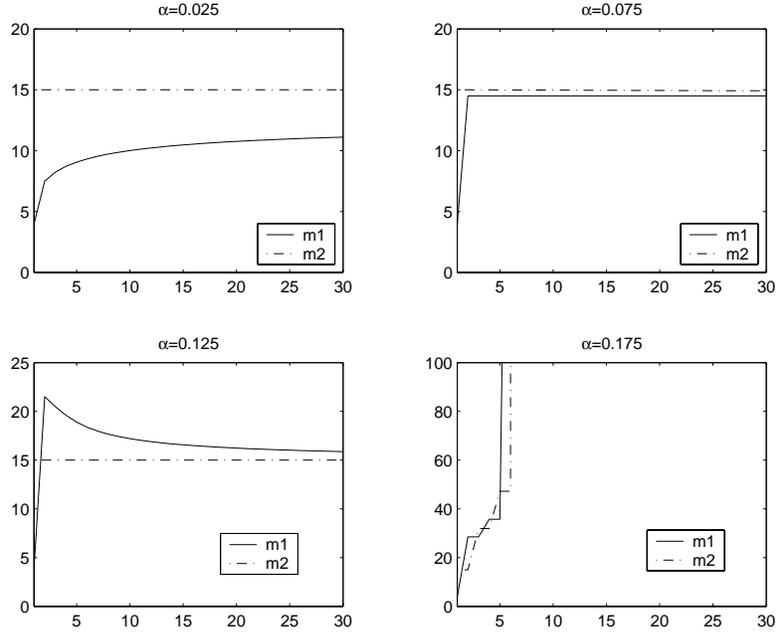


Figure 7: Firm dynamics for different values of α . $n = 40$.

$m_1(\mathbf{0}) > m_2(\mathbf{0})$ and $\mathbf{0}$ otherwise, was included since it was found to be statistically significant.

Column (4) of table 1 also includes coefficients for the normalized variables, which all have mean of zero and standard deviation of one. These coefficients then allow us to compare the relative magnitudes of the variables on the dependant variable.

>From the regression table we can draw the following conclusions:

- Increasing initial own firm size, increases long run size, and increasing rival's size initially has a negative effect, but for $m_2(\mathbf{0}) > 8$, the effect is positive.
- The parameters α and β both have positive effects on firm size, though negative interaction effects with initial size.
- The dummy variable, which reflects the relative starting position of the two firms, shows that relative initial size matters in determining long

Dependent Variable: $m_1(100)$			
Variable	Coef.	Std. Err.	Norm. Coef.
$m_1(0)$	0.517	0.032	0.475
$[m_1(0)]^2$	0.009	0.001	0.190
$m_2(0)$	-0.289	0.030	-0.264
$[m_2(0)]^2$	0.031	0.001	0.632
$m_1(0) \cdot m_2(0)$	-0.020	0.001	-0.308
$\delta [m_1(0) > m_2(0)]$	0.447	0.090	0.038
α	58.159	3.763	0.208
β	159.342	4.536	0.569
$\alpha \cdot m_1(0)$	-3.110	0.310	-0.175
$\beta \cdot m_1(0)$	-2.042	0.289	-0.114
$\beta \cdot m_2(0)$	3.907	0.309	0.216
n	-0.040	0.008	-0.077
$n \cdot \beta$	1.401	0.125	0.163
$n \cdot m_1(0)$	0.006	0.001	0.170
<i>Constant</i>	3.157	0.301	.
Nobs.	5000		
R^2	0.878		

Table 1: Regression results for adjustment dynamics. Dep. Var. $m_1(100)$. Robust standard errors given. All variables stat. sig. at 99% or greater confidence level.

run size. If firm 1 starts larger than firm two, all else equal, it will have a larger long run size.

- **Ceteris paribus**, increasing complexity is associated with smaller firm size. The reason is that greater environmental complexity reduces profits and thus in turn reduces the long run size.
- However, there are several interaction effects that capture the non-linear relationship between the independent variables and long run firm size. For example, both initial own size and the adjustment parameters have positive effects, but there are also off-setting interaction effects that reduce the total effect of starting size and the adjustment parameters. That is to say, the interaction variables, $\alpha \cdot m_1(0)$ and $\beta \cdot m_1(0)$ have negative effects, while $n \cdot \beta$ has a positive effect. Interestingly,

$n \cdot \alpha$ has no statistically significant effect on long run size (coefficient not included).

7 Convergence to Nash Equilibrium

As we saw in the previous section, the long run size of the firm is determined by several variables and, as a result, the convergence to the Nash equilibrium is not guaranteed by the simple adaptive dynamics that we study in this paper. As a result, this section investigates the conditions under which convergence takes place.

We made random draws of the relevant parameters (α and β in the range $[0, 0.2]$, $m_1(0)$ and $m_2(0)$ in the range $[2, 20]$), and we ran the dynamics. Then, we retained only the runs that ended up with both firms within one node of the Nash value (i.e. $m_i(50) \in [m^* - 0.5, m^* + 0.5]$). That was done one million times for each complexity value $n \in \{5, 10, 15, 20, 25, 30, 35, 40, 45\}$.

Table 2 shows the results, reporting the success rate, the average α , β , and initial m_1 , and the mode of the latter (given the symmetric setting, the values for $m_2(0)$ are identical). The numbers in parenthesis are the standard deviations.

n	(m^*)	$Succ.$	$\bar{\alpha}$	$\bar{\beta}$	$m_1(0)$	$mod[m_1(0)]$
5	(6.85)	0.492%	0.122 (0.052)	0.082 (0.042)	8.758 (5.380)	5
10	(7.19)	0.413%	0.115 (0.057)	0.087 (0.056)	7.897 (5.267)	4
15	(7.58)	0.296%	0.093 (0.056)	0.082 (0.075)	4.137 (1.476)	5
20	(8.01)	0.190%	0.106 (0.053)	0.009 (0.007)	4.845 (1.792)	5
25	(8.53)	0.207%	0.106 (0.054)	0.011 (0.007)	4.972 (1.896)	5
30	(9.15)	0.247%	0.105 (0.053)	0.012 (0.008)	5.408 (2.137)	5
35	(9.91)	0.293%	0.105 (0.053)	0.014 (0.010)	5.700 (2.347)	6
40	(10.93)	0.354%	0.110 (0.050)	0.016 (0.011)	6.304 (2.583)	5
45	(12.50)	0.356%	0.119 (0.048)	0.021 (0.014)	6.967 (2.948)	6

Table 2: Convergence to Nash equilibrium for different complexity levels

The first result that emerges is that only a very small number of runs converged close to the Nash value. In fact, the success ratio does not even attain half of a percentage point. The value was particularly low for intermediate complexity values. Thus, we find that even extremely simple and commonsensical adjustment rules, while allowing for convergence to a steady state value, do not yield the equilibrium 'full information' outcome. This result calls for a careful assessment of the conditions under which a Nash outcome may be seen as a plausible outcome.

Another striking result is that the mode of the initial size is in fact insensitive to complexity, whereas the mean has a slight u-shape (i.e., larger values at the extremes of the complexity range). Thus, with respect to both statistics, we do not observe that the increase in complexity, and in the associated Nash value for firm size, requires larger initial firms for convergence to happen.

In fact, from the table we can induce that the only variable that seems to vary significantly with complexity is the profit coefficient β . The others are quite similar across complexity values, both regarding the value and the standard deviation.

To further investigate the relationship between the exogenous variables and the likelihood of the system settling down to a Nash equilibrium, we conducted a probit regression analysis. We created data set of **50,000** observations from random draws of the relevant parameters from the respective parameter sets given above. The probit analysis allows us to measure how the parameters affect the probability that the system will reach a Nash equilibrium. The results of this regression are given in table 3, where we report the results of the marginal changes in probability with a change in the independent variable. A few conclusions can be drawn from the table:

- The probability of firm size reaching a Nash equilibrium is decreasing with environmental complexity. This is quite intuitive, as more complex environments have a perturbing effect on the dynamics
- **Cet. par.** larger α is associated with a higher probability of reaching Nash. This is due to the fact that the α parameter reflects the firm interaction with the opponent, and thus may be a proxy for a best response dynamics. The more important this factor in the adjustment dynamics, the greater the chances to mimic the best response dynamics.

Variable	dF/ dx	std. err.
n	-0.00015	3.51E-05
n^2	2.90E-06	6.55E-07
α	0.0046	0.0014
β	-0.0583	0.0072
β^2	0.170	0.028
$m_1(\mathbf{0})$	-0.00039	5.66E-05
$m_2(\mathbf{0})$	-0.00039	5.58E-05
$\beta \cdot m_1(\mathbf{0})$	0.00089	0.00024
$\beta \cdot m_2(\mathbf{0})$	0.00047	0.00019
$m_1(\mathbf{0}) \cdot m_2(\mathbf{0})$	0.000022	3.65E-06
nobs.	50,000	
prob(Nash)	0.0033	
prob(Nash)	.00076	(at \bar{x})
pseudo R^2	0.243	

Table 3: Probit regressions results. Dep. var. is **1** if long run firm size is equal to Nash equilibrium, **0** otherwise. Robust standard error are presented. All coefficients are stat. sig. at 99% or greater confidence level.

- β has a quadratic relationship: having a negative effect for small betas but increasing for values of β above **0.325**.
- Initial firm size matters; the larger the initial size the lower the probability of hitting the Nash equilibrium; though there are positive, off-setting interaction effects between own and rivals sizes.
- Lastly there is a positive interaction effect between β and initial own and rival's size.

8 Discussion and Conclusion

This paper has presented a model of the firm as an artificial neural network. We explored the long run size dynamics of firms/neural networks playing a Cournot game in an uncertain environment. Building on previous papers we derived a profit function for the firm that depends on its own size, the rival's size and environmental complexity. We then looked at long-run firm

size resulting from two types of simple adaptive rules: the 'isolationist' and the 'imitationist.' These dynamics were compared to a benchmark 'best response' case.

First we find that when using simple adjustment rules, long run firm size is a function of initial firm size, initial rival's size, environmental complexity and the adjustment rate parameters. These variables interact in a non-linear way. We also find that only under very precise initial conditions and parameter values does the firm converge to the Nash equilibrium size given by the best response. The reason is that the dynamics we consider tend to settle rapidly (no more than a few iterations) on a path that depends on initial conditions. The simple rules generally yield suboptimal long run equilibria, and only fairly specific combinations of parameters and initial sizes yield the optimal size. Thus, our results suggest cautiousness in taking the Nash equilibrium as a focal point of simple dynamics (the standard 'as if' argument). Interestingly, we further find that when firms use simple adjustment rules, environmental complexity has a negative effect on size, even if the Nash equilibrium size is increasing in environmental complexity. This is particularly true in the isolationist case, and is explained by the negative correlation between profits and complexity: complex environments yield lower profits, and hence less incentives for firm growth; *ceteris paribus*, this yields lower steady state size, and suboptimal profits. Thus, in our model more efficient information processing, and more complex adjustment rules would play a positive role in the long run profitability of the firm, and would deserve investment of resources. It may be interesting to build a model in which firms face a trade-off between costs and benefits of costly but efficient adjustment rules.

Finally, we found that β , the effect of own profit on firm size, is more important than the comparison with the rival's performance in determining whether the Nash outcome is reached or not. In fact, it emerged from our analysis that values of β in particular ranges significantly increased the probability of convergence to the optimal size. This result triggers the question of why such a parameter should be considered exogenous. One could imagine a 'superdynamics', in which firms adjust their reactivity to market signals, i.e., their β , in order to maximize the chances of converging to the optimal size. Such an investigation may also be the subject of future research. Other extensions may include to link the findings with these models to stylized facts around firm growth and size (i.e., the literature on Gibrat's Law).

Appendix

A Neural Networks

This appendix briefly describes the working of Artificial Neural Networks.. For a more detailed treatment, the reader is referred to Skapura (1996). Neural networks are nonlinear function approximators that can map virtually any function. Their flexibility makes them powerful tools for pattern recognition, classification, and forecasting. The Backward Propagation Network (BPN), which we used in our simulations, is the most popular network architecture. It consists in a vector of $\mathbf{x} \in \mathbf{R}^n$ inputs, and a collection of m processing nodes organized in layers.

For our purposes (and for the simulations) we focus on a network with a single layer. Inputs and nodes are connected by weights, $\mathbf{w}^h \in \mathbf{R}^{n \times m}$, that store the knowledge of the network. The nodes are also connected to an output vector $\mathbf{y} \in \mathbf{R}^o$, where o is the number of outputs (2 in our case), and $\mathbf{w}^o \in \mathbf{R}^{m \times o}$ is the weight vector. The learning process takes the form of successive adjustments of the weights, with the objective of minimizing a (squared) error term.⁵ Inputs are passed through the neural network to determine an output; this happens through transfer (or squashing) functions, such as the sigmoid, to allow for nonlinear transformations. Then supervised learning takes place in the sense that at each iteration the network output is compared with a known correct answer, and weights are adjusted in the direction that reduces the error (the so called ‘gradient descent method’). The learning process is stopped once a threshold level for the error has been attained, or a fixed number of iterations has elapsed. Thus, the working of a network (with one hidden layer) may be summarized as follows. The feed forward phase is given by

$$\hat{\mathbf{y}}_{1 \times o} = g \left(g \left(\mathbf{x}_{1 \times n} \cdot \mathbf{w}^h_{n \times m} \right) \cdot \mathbf{w}^o_{m \times o} \right)$$

where $g(\cdot)$ is the sigmoid function that is applied both to the input to the hidden layer and to the output. To summarize, the neural network is comprised of three ‘layers’: the environmental data (i.e., the environmental state vectors), a hidden/managerial layer, and an output/decision layer. The ‘nodes’

⁵The network may be seen as a (nonlinear) regression model. The inputs are the independent variables, the outputs are the dependent variables, and the weights are equivalent to the regression coefficients.

in the managerial and decision layers represent the information processing behavior of agents in the organization.

The error vector associated with the outputs of the network is:

$$\varepsilon = \sum_{j=1}^o (y_j - \hat{y}_j)^2, \quad j = 1, \dots, o$$

Total error is then calculated:

$$\xi = \sum_{j=1}^o \varepsilon_j$$

where \mathbf{y} is the true value of the function, corresponding to the input vector \mathbf{x} .

This information is then propagated backwards as the weights are adjusted according to the learning algorithm, that aims at minimizing the total error, ξ . The gradient of ξ with respect to the output-layer weights is

$$\frac{\partial \xi}{\partial \mathbf{w}^o} = -2(y - \hat{y}) [\hat{y}(1 - \hat{y})] g'(\mathbf{x} \cdot \mathbf{w}^o),$$

since for the sigmoid function, $\partial \hat{y} / \partial \mathbf{w}^o = \hat{y}(1 - \hat{y})$.

Similarly, we can find the gradient of the error surface with respect to the hidden layer weights:

$$\frac{\partial \xi}{\partial \mathbf{w}^h} = -2 \sum_{j=1}^o (y_j - \hat{y}_j) [\hat{y}_j(1 - \hat{y}_j)] g'(\mathbf{x} \cdot \mathbf{w}^h) \mathbf{w}^o g'(\mathbf{x} \cdot \mathbf{w}^o) \mathbf{x}.$$

Once the gradients are calculated, the weights are adjusted a small amount in the opposite (negative) direction of the gradient. We introduce a proportionality constant η , the learning-rate parameter, to smooth the updating process. Define $\delta^o = .5(y - \hat{y}) [\hat{y}(1 - \hat{y})]$. We then have the weight adjustment for the output layer as

$$\mathbf{w}^o(t + 1) = \mathbf{w}^o(t) + \eta \delta^o g'(\mathbf{x} \cdot \mathbf{w}^o)$$

Similarly, for the hidden layer,

$$\mathbf{w}^h(t + 1) = \mathbf{w}^h(t) + \eta \delta^h \mathbf{x},$$

where $\delta^h = g'(\mathbf{x} \cdot \mathbf{w}^h) \delta^o \mathbf{w}^o$. When the updating of weights is finished, the firm views the next input pattern and repeats the weight-update process.

B Derivation of the Profit Function

This appendix briefly describes the process that leads to equation 2, that in the present paper is left in the shadow. Details of the model can be found in BS (2004) and BS (2005).

B.1 Cournot Competition in an Uncertain Environment

We have two Cournot duopolists facing the demand function

$$p_t = \gamma_t - (q_{1t} + q_{2t}).$$

where γ_t changes and is **ex ante** unknown to firms. Assume that production costs are zero. Then, the best response function, were γ_t known, would be given by

$$q_j^{br} = \frac{1}{2}[\gamma - q_{-j}],$$

with a Nash Equilibrium of

$$q^{ne} = \frac{\gamma}{3}, \quad \pi_j^{ne} = \frac{1}{9}\gamma^2.$$

When deciding output, firms do not know γ , but have to estimate it. They only know that it depends on a set of observable environmental variables $\mathbf{x} \in \{0, 1\}^n$:

$$\gamma_t = \gamma(\mathbf{x}_t) = \frac{1}{2^n} \sum_{k=1}^n x_k 2^{n-k}$$

where $\gamma(\cdot)$ is unknown **ex ante**. Each period, the firm views an environmental vector \mathbf{x} and uses this information to estimate the value of $\gamma(\mathbf{x})$. Note that $\gamma(\mathbf{x}_t)$ can be interpreted as a weighted sum of the presence or absence of environmental features.

To measure the complexity of the information processing problem, we define environmental complexity as the number of bits in the vector, n , which, ranges from a minimum of **5** bits to a maximum of **50**. Thus, in each period:

1. Each firm observes a randomly chosen environmental state vector \mathbf{x} . Note that each \mathbf{x} has a probability of $1/2^n$ of being selected as the current state.

2. Based on that each firm estimates a value of the intercept parameter, $\hat{\gamma}_j$. The firm also estimates its rival's choice of output, \hat{q}_{-j}^j , where \hat{q}_{-j}^j is firm j 's guess of firm $-j$'s output.
3. It then observes the true value of and γ , and q_{-j} , and uses this information to determine its errors using the following rules:

$$\varepsilon_{1j} = \hat{\gamma}_j - \gamma \quad \phi_2 \quad (6)$$

$$\varepsilon_{2j} = \hat{q}_{-j}^j - q_{-j} \quad \phi_2 \quad (7)$$

4. Based on these errors, the firm updates the weight values in its network.

This process repeats for a number $T = 250$ of iterations. At the end, we can compute the average profit for the two firms as

$$\pi_i = \frac{1}{T} \sum_{t=1}^T q_{it}(\gamma_t - (q_{1t} + q_{2t})). \quad (8)$$

C Regression Results for Profit

Equation (2) was derived by using the model described in the preceding appendix. We built a data set by making random draws of $n \in [5, 50]$, $m_i \in [2, 20]$. We ran the Cournot competition process for $T = 250$ iterations (random initial conditions were appropriately taken care of by averaging over multiple runs). We recorded average profit for the two firms computed as in eq. (8), and the values of m_1 , m_2 , and n . This was repeated 10,000 times, in order to obtain a large data set. We then ran a regression to obtain a precise polynomial form for profit as a function of sizes and environmental complexity. Table 4 gives the complete results of the regression, which is reflected in equation (1).

Variable	Coefficient	Std. Error
<i>constant</i>	270.688	0.898
m_1	5.932	0.229
m_1^2	-0.375	0.023
m_1^3	0.007	0.001
m_2	0.490	0.056
$m_1 \cdot m_2$	-0.304	0.004
n	-2.201	0.034
n^2	0.003	0.001
$m_1 \cdot m_2 \cdot n$	0.007	0.000
$m_2 \cdot n$	-0.016	0.002
nobs.	10,000	
R^2	0.864	
\bar{R}^2	0.864	

Table 4: Profit function for firm 1. Dep. var $10,000 \cdot \pi_1$. Robust standard errors given. All coefficients stat. sig. at 99% or greater confidence level.

References

- Barr, J. and Saraceno, F. (2005). "Cournot Competition, Organization and Learning." *Journal of Economic Dynamics and Control*, 29(1-2), 277-295.
- Barr, J. and Saraceno, F. (2004). "Organization, Learning and Cooperation." Rutgers University, Newark Working Paper #2004-001.
- Barr, J. and F. Saraceno (2002), "A Computational Theory of the Firm," *Journal of Economic Behavior and Organization*, 49, 345-361.
- Chandler, Jr., A. D. (1977). *The Visible Hand: The Managerial Revolution in American Business*. Harvard University Press: Boston.
- Chang, M-H, and Harrington, J.E. (forthcoming). "Agent-Based Models of Organizations." *Handbook of Computational Economics*, Vol 2. Eds K.L. Judd and L. Tesfatsion.
- Cyert, R. M. and March J.G. (1963). *A Behavioral Theory of the Firm*. Prentice-Hall, New Jersey.

- DeCanio, S. J. and W. E. Watkins (1998), "Information Processing and Organizational Structure," **Journal of Economic Behavior and Organization**, 36, 275-294.
- Fudenberg, D. and Tirole, J. (1996). **Game Theory**. The MIT Press, Cambridge.
- Nelson, R. R. and Winter, S. G. (1982). **An Evolutionary Theory of Economic Change**. Belknap Press of Harvard University Press: Cambridge.
- Radner, R. (1993), "The Organization of Decentralized Information Processing," **Econometrica**, 61, 1109-1146.
- Simon, H. A. (1982). "Rational Choice and the Structure of the Environment." **Psychology Review**. 63(2): 129-138. Reprinted in: **Behavioral Economics and Business Organizations**. The MIT Press: Cambridge.
- Skapura, D. M. (1996), **Building Neural Networks**, Addison-Wesley, New York.
- Van Zandt, T. (1998). "Organizations with an Endogenous Number of Information Processing Agents," In: M. Majumdar (Ed.) **Organizations with Incomplete Information**. Cambridge University Press: Cambridge.