Organization, Learning and Cooperation*

Jason Barr†and Francesco Saraceno‡
Rutgers University, Newark Working Paper #2004-001

March 2004

Abstract

We model the organization of the firm as a type of artificial neural network in a duopoly framework. The firm plays a repeated Prisoner's Dilemma type game, but also must learn to map environmental signals to demand parameters. We study the prospects for cooperation given the need for the firm to learn the environment and its rival's output. We show how a firm's profit and cooperation rates are affected by its size, its rival's size and willingness to cooperate, and environmental complexity.

JEL Classification: C63, C72, D21, D83, L13
Keywords: Artificial Neural Networks, Cooperation, Firm Learning

*An earlier version of this paper was been presented at the 2003 Annual Meeting of the Society for the Advancement of Behavioral Economics, Lake Tahoe, Nevada, July 29-31. We thank Duncan Foley and seminar participants at the U.S. Naval Academy for their comments and suggestions. The usual caveats apply.

†Corresponding Author: Department of Economics, Rutgers University, Newark, NJ 07102. email: jmbarr@andromeda.rutgers.edu.

‡Observatoire Français des Conjonctures Économiques, Paris, France. email: francesco.saraceno@sciences-po.fr

1 Introduction

Information processing and decision making by firms are not typically done by one person. Rather decisions are made by a group of people either in a committee or hierarchical structure. Bounded rationality and/or computational costs preclude the possibility of any one agent collecting, processing and deciding about information relevant to the firm and its profitability. Thus many agents are employed to process this information so that the firm can make informed decisions. But processing information and decision making are costly activities. Large firms, for example, employ hundreds even thousands of 'managers' who do not produce or sell anything, but rather process information and make decisions (Radner, 1993).

Building on previous work (Barr and Saraceno, forthcoming), we model the firm as a type of artificial neural network (ANN) that must make output decisions in a Cournot duopoly framework. The use of ANNs as a model of firm organization allows us to make explicit the nature and cost and benefits of processing information. Agents within the firm are required to evaluate data and communicate this evaluation to others who then make final decisions. As discussed in Barr and Saraceno (2002; forthcoming), the benefit to the firm of increased resources devoted to information processing (IP) is better knowledge of the environment and hence increased returns. But the costs include the costs of paying agents and the time and costs involved with processing and communicating this information.

In Barr and Saraceno (forthcoming) we investigated how learning affects a firm's ability to produce along the best response function. We show that ANNs competing in a duopoly setting can learn to converge to the Nash equilibrium output, which is changing each period due to changing environmental states. Further we show how environmental complexity affects both organizational size and profitability. In that setting, firms choose an output level given the observation of the environmental state and, after that, the price is set, the market clears, and firms then compare their output choices to the best response they should have played given the output choice of their opponent.

The ANN embodies the knowledge of the firm and implicitly contains a history of what the rival played in the past. Over time, firms learn to forecast both the effect of changing environmental states on demand and the rival's output; both firms converge to the Nash equilibrium over time. In that framework, we did not consider the possibility of collusion in the sense that firms could possibly learn to produce less than Cournot output, thus gaining larger profits.

Here, we are interested in studying the prospects for cooperation with two additional factors as compared to traditional cooperation models (see Axelrod, 1984): (1) the need for firms to map environmental characteristics to changing product demand, and (2) many agents within a firm are needed to learn not only the environment but what exactly the rival is signalling about its output strategy, i.e., whether it will cooperate or defect. We argue that a firm's knowledge not only of the environment but also its ability to see what its rival is doing will ultimately affect its profitability. If, for example, a firm learns that its rival has a low desire to cooperate, all else equal, it too should change its output to defect most of the time. But if a firm learns that its rival is willing to cooperate then mutual gains are possible. The difficulty of learning affects a firm's ability to cooperate, since in highly complex environments, firms will have problems distinguishing what part of the variation in the market clearing price is coming from the variation in the environment versus the variation in the rival's strategy.

In this paper, we have two neural networks competing in a repeated duopoly setting. Each period, the firm (network) views an environmental state vector (an N-length vector of $0's$ and $1's$) which it uses to estimate two variables: the intercept parameter of the demand curve, and its rival's output. The firm uses the estimate of its rival's output to decide if its rival is going to defect or cooperate. If the estimate of its rival's output is greater than the shared-monopoly output (plus some margin), the firm plays the (estimated) Cournot best response, otherwise it plays the (estimated) shared-monopoly response. In this sense, both firms are playing a kind of Tit-For-Tat strategy: 'I defect if I think you will defect, otherwise I will cooperate.'

After each firm chooses an output, it observes the market clearing price, the rival's output and the true demand intercept. The firm uses this information to calculate the error it made in its estimation of the intercept and its rival's output, and then uses this error to improve its performance in the subsequent periods. Thus the knowledge of both the environmental characteristics and the rival's past plays lies within the network itself, rather than any one agent; the network serves as a economical storage device: 'given what I have learned in the past, I now can map environmental information to a rival's output and to the demand intercept.'

We measure **environmental complexity** by the number of environmental bits (signals) a firm views each period. Each period a new environmental vector is randomly chosen with probability of $1/2^N$, where N is the number of bits in the vector. This is tantamount to what might be referred to as a 'pure generalization' process.[1] In our paper, for N large enough, this has the effect of having an environment that is changing at each period.

In this paper we ask:

- What is the relationship between network size, learning and profitability given that firms are learning both the environment and their opponent's behavior?
- How does environmental complexity affect performance, cooperation and profits?
- What is the relationship between a firm's willingness to be 'nice,' i.e., its willingness not to defect given that it estimates its rival will defect, profits and cooperation?
- What is the average firm size and cooperation in equilibrium versus environmental complexity?

To anticipate some of our results, we will be drawing three sets of conclusions. The first is that environmental complexity and firm dimension interact in a complex way: performance is a function of the not only the difficulty of the IP task, but also the rival's size and willingness to cooperate. Second, we will show that cooperation rates are also affected by environmental complexity, firm size, and rival's behavior. Lastly, we will show how strategic interaction leads to 'industrial equilibria' in regards to firm size and cooperation; and we show how environmental complexity affects these equilibria.

The rest of the paper is as follows. In the next section we briefly review the literature that relates to our paper. The following (section 3) outlines the standard Cournot model that we work with, which is an extension of the RPD

[1] In the computer science literature, the type of ANN that we employ–the Backward Propagation Network–normally 'trains' on a fixed data set and then is presented new data for forecasting (Croall and Mason, 1992). After training the network can 'generalize' in the sense that it can make forecasts using new, unseen data. Here, the networks train and generalize simultaneously.

game. Next, section 4 discusses the particular game that the neural networks play and the characterization of the economic environment. In section 5 we discuss the workings of the particular ANN that we use. Then sections 6 through 8 present the results of our simulation experiments. Finally, in section 9 we present concluding remarks and possible research extensions.

2 Related literature

There is a rich literature on the issue of cooperation and defection in Prisoner's Dilemma type games, of which the Cournot game is a variant. In this general framework, two players must make a decision over two outcomes: whether to cooperate or defect. The result (payoff) of the decision, however, is affected by the rival's decision. If decisions have to be made repeatedly and there is little prospect that the game will end in a short time then mutual sustained cooperation is possible, even if there is no direct communication. In the Cournot game, firms can signal their willingness to cooperate over time by choosing low output over high output. If the other firm 'takes the bait' by also producing a low output, then their is mutual gain to be had (assuming a sufficiently high discount rate).

In terms of the repeated Prisoner's Dilemma (RPD), a standard game theoretic result is that of the 'folk theorem,' which says that if agents are patient enough, then there are an infinite set of Nash Equilibrium outcomes that have higher pay-offs than the 'defect every period' strategy (the so-called min-max payoff) (Fudenberg and Tirole, 1991).

Recently, models of the RPD have also been concerned with bounded rationality and the evolution of cooperation. Rubinstein (1986) and Cho (1994) model agents as boundedly rational automata-type machines. Rubinstein's machine is a finite automata and Cho's is a simple perceptron. These papers show the types of equilibria that can arise. If, for example, there is a bound placed on level of complexity in Rubinstein's machine then only a finite number of equilibrium outcomes can be generated. Cho's machine is able to 'recover' the perfect folk theorem using a neural network that maintains an upper bound on the complexity of equilibrium strategies.

While these papers focus on the nature of the machines and the nature of the equilibria outcomes, other papers focus on the evolution of cooperation (Axelrod, 1997; Miller, 1996; Ho, 1996). For example, Miller demonstrates how cooperation can evolve overtime if automata machines adapt using a genetic algorithm (Holland, 1975) that allows the strategic environment to change. Further he studies co-evolution of strategies under imperfect information. Axelrod (1997) also models the RPD with a genetic algorithm, but fixes the population of possible strategies.

Our paper relates to the literature on adaptive machines that play a repeated Prisoner's Dilemma type game but is different in the following ways. First, we are interested in studying an RPD using an agent-based model of the firm. Our interest is in asking the questions: How are profits and cooperation affected by agent-based learning? And what are the optimal number of agents needed to learn both the external economic environment and the rival's output decision over time? For simplicity we hold the firms' strategies constant (as a type of Tit-for-Tat strategy) and focus on the relationship between firm complexity (network size), environmental complexity (the quantity of information), profitability and cooperation. Thus our objective is to focus not on the evolution of strategies or the types of equilibria outcomes but rather the learning process that firms need to do in order to improve performance and profits.

In addition to RPD games, cooperation in Cournot models have been widely discussed (see Tirole (1988) for a review of these models). Similar to the RPD, collusion is possible if firms are sufficiently patient and the threat of punishment exists (Verboven, 1997). Cyret and DeGroot (1973) show cooperation is possible if firms maximize joint profits; further they can come to cooperate over time by a process of Bayesian learning. Vriend (2000) presents an adaptive model of a Cournot game, where agents evolve according to a genetic algorithm. He shows how equilibrium market outcomes can be different depending if agents perform individual rather than social learning. Our model is similar to these papers in the sense that firms are adaptive, but unlike these models, firms's output decisions evolve based on not only the rival's behavior but also the nature of the environment and on the inherent 'niceness' of firms themselves.

Our work also fits within the literature on agent-based models of the firm (Radner, 1993; DeCanio and Watkins, 1998; Carley, 1996). These models, borrowing heavily from computer science, represent the firm as a network of information processing agents (nodes). In general these papers study which types of networks minimize the costs of processing and communicating information. Our model is also agent-based in the sense that we assume that output decisions by the firm are made by a network of information processing agents. However, our work is different in two respects: in general, and unlike other agent-based models, we directly model the relationship between the external environmental variables, firm learning and performance; secondly, we explicitly provide an agent-based model of Cournot competition and cooperation, which to our knowledge has not been done before.

Our agent-based approach models the firm as a type of artificial neural network, with the nodes representing managers. ANNs are common in computer science and psychology, where they have been used for pattern recognition and modeling of the brain (Croall and Mason, 1992; Skapura, 1996). In economics, neural networks have been employed less frequently. One application has been to use ANNs as non-linear estimation equations (Kuan and White, 1992). Because of the stochastic and non-linear nature of ANNs we employ a simulation-based approach to studying the relationship between firm performance, competition and size.

3 The duopoly framework

This section will give a textbook summary of standard Cournot theory in a static and repeated framework. Then, in the next section we will introduce uncertainty and show how we model firms as ANNs.

Let's say we have a market with two firms. Each period they face the demand function

$$p_t = \alpha_t - \beta \left(q_{1t} + q_{2t} \right).$$

In section 4, α will be a variable, and we will assume that firms must estimate its value from period to period; but for the moment, to review the Cournot game, we take it as constant and known to the firms. Here and for the rest of the paper, we also assume that the slope is constant (and normalized to 1). Profits for each firm are

$$\pi_j = \left[\alpha - (q_1 + q_2) \right] q_j - c_j, \ j = 1, 2$$

where c_j is costs, such as the cost of network, and is set to zero for convenience and without loss of generality.[2]

[2] While we do not deny the importance of the cost of carrying a network of a given size, we do not include this cost in the paper for simplicity since the qualitative results would not change.

Under standard Cournot assumptions, the best response function is given by

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

with a Nash Equilibrium of

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

This is a typical prisoner dilemma's game. If the two firms could coordinate their output decisions and act as a monopoly their joint profit from production would be

$$\pi^m = \left[\alpha - Q\right]Q.$$

Assuming they share production and profits equally, each would have an optimal output:

$$q_j^m = \frac{\alpha}{4}.$$

Profit would then be

$$\pi_j^m = \frac{1}{8}\alpha^2 > \pi_j^{ne} = \frac{1}{9}\alpha^2.$$

In a single shot game the cooperation outcome is not an equilibrium;[3] if one firm knew that its rival would play half of the monopoly output, then it could defect by playing its Cournot best response and achieve a higher payoff:

$$\begin{aligned} q_j^d &= \frac{1}{2}\left[\alpha - q^m\right] = \frac{3}{8}\alpha \\ \pi_j^d &= \frac{9}{64}\alpha^2 > \pi_j^m. \end{aligned}$$

On the other hand, as is well known, if the game is repeated other equilibrium outcomes can emerge.

[3]Here we define cooperation as splitting the monopoly quantity and defection is each playing a best response.

3.1 Repeated prisoner dilemmas and the folk theorem

In a repeated framework, firms face the decision each period wether to 'co-operate' or 'defect.' Then, whether the horizon is finite or not will yield completely different results. Suppose that the game is repeated an infinite number of times, and that firm j's payoff is given by the discounted sum of profits:

$$\Pi_j = \sum_{t=1}^{\infty} \delta^{t-1} \pi_{jt}$$

where δ is the discount rate.[4] In this case many other strategies involving cooperation can constitute an equilibrium, in addition to the 'defect every period' outcome. Take for example the grim trigger strategy: cooperate until the opponent defects; if defection revert to the best response strategy (i.e., defect) forever. The cooperative outcome will be assured when

$$\frac{\pi^m}{1-\delta} > \pi^d + \pi^{ne} \frac{\delta}{1-\delta}.$$

Plugging in the profit values from above, we have for our case $\delta^* > \frac{9}{17}$. In words, if the discount rate is large enough, future losses will more than compensate the short run gain from defection, and cooperation will be the outcome of a grim trigger strategy. But the same could be showed for the so called Tit-For-Tat strategy, which consists of beginning with cooperation, and playing each period what the opponent played in the preceding one. In fact, it can be shown that if the discount rate is high enough, almost any strategy involving cooperation can be an equilibrium strategy. This is in fact what the folk theorem tells us: any feasible expected payoff can be sustained in an equilibrium as long as each player can expect an equilibrium payoff larger than the uncooperative one. In that case, no player will have an incentive to deviate. With a backward induction argument, on the other hand, it can be shown that cooperation is not sustainable if the game is repeated a finite number of times.

An important extension of the preceding framework that relates to the present paper is the consideration of the effects of uncertainty on the sustainability of cooperative equilibria (Green and Porter, 1984). This is typically

[4]Notice that δ can either be interpreted as the discount factor of an infinitely repeated game, or as the probability that the game is repeated after each round when the game length is undefined. analytically the two 'stories' boil down to the same thing.

modelled as price uncertainty; for example, the constant term of a linear demand function shifts according to a given probability distribution. The obvious effect of such a feature of the model is that deviations from the collusive price and profit are not directly attributable to the competitor's unwillingness to cooperate, but may stem from shifts in the demand function. The punishment scheme designed by a firm to force cooperation has as a consequence to be more complex than in the case of certainty. This typically involves a trade-off, whose outcome depends on the particular model adopted: if the punishment is too harsh, the firm loses possible advantages from collusion; but if it is too light, then the opponent may be tempted to adopt a noncooperative stance. To sustain cooperation, hence, firms have to punish their opponents only if prices and profits deviate "too much" from the cooperative level. We'll adopt a similar perspective in what follows, with a crucial difference. We add the possibility that firms can reduce uncertainty by means of learning; our focus, in fact, is on this learning process, and on how it affects the willingness to cooperate.

4 A model of firm learning in a repeated Cournot game

4.1 Strategies

In this paper we have each network employ a relatively simple strategy: a type of Tit-for-Tat (TFT). The standard TFT says that a firm should begin by cooperating and then play the same outcome as the rival's prior move. Given our framework, this strategy allows for the possibility of cooperation once firms begin to learn the external environment, and learn to separate the variance in price that is due to environmental change versus their rival's output decision.[5]

In this paper, firms employ a slight variation of the TFT strategy. Since

[5] TFT, in general, behaves according to the four rules-of-thumb discussed by Axelrod (1984) for strategies that are likely to promote cooperation in a setting where boundedly rational agents are the players: (1) Be nice: never be the first to defect. (2) Be forgiving: be willing to return to cooperating even if your opponent defects. (3) Be simple: the easier it is to discover a pattern in a rival's output the easier it is to learn to cooperate. (4) Don't be envious: don't ask how well you are doing compared to your rival, but rather how much better you can do, given your rival's actions.

they estimate both the demand parameter and their rival's output quantity each period, they have to use this information to decide whether to defect or not. More specifically the firm chooses an output each period based on the following rule ($j = 1, 2,$):

$$q_j = \left\{ \begin{array}{c} \frac{1}{2}\left(\hat{\alpha}_j - \hat{q}^j_{-j}\right) \text{ if } \left(\hat{q}^j_{-j} - \frac{\hat{\alpha}_j}{4}\right) > \rho_j \\ \hat{\alpha}_j/4 \text{ otherwise} \end{array} \right\}, \tag{1}$$

where q_j is firm $j's$ output, $\hat{q}^j_{-j}$ is firm $j's$ estimate of its rival's output, and $\hat{\alpha}_j$ is firm $j's$ estimate of α. Equation (1) says that if the firm estimates its rival to be a cheater: $\hat{q}^j_{-j} > \hat{\alpha}_j/4 + \rho_j$, i.e., that the rival is expected to deviate from forecast monopoly profit, then it plays the optimal forecasted Cournot output; that is, it defects as well.

The threshold value $\rho_j \geq 0$ represents the firm's 'willingness to be nice.' For relatively small values, e.g., $\rho_j = 0$, firm j will play defect relatively more often; for values $\rho_j \geq \ddot{\rho}$, the firm will be so nice that it will never defect. Notice that in making this decision the firm has two possible sources of error: the first is the environment, and the second is the opponent's quantity; this is why it will allow a deviation ρ_j from the monopoly output before reverting to the noncooperative quantity.

4.2 The economic environment

4.2.1 A shifting demand curve

Here we represent the external environment as a vector of binary digits $\mathbf{x} \in \{0, 1\}^N$. The relationship between the environment and the intercept is given by

$$\alpha(\mathbf{x}) = \frac{1}{2^N} \sum_{k=1}^{N} x_k 2^{N-k},$$

x_i is the k^{th} element of $\mathbf{x}$. This functional relationship converts a binary digit vector into its decimal equivalent.[6] We can think of this in the following manner: the vector $\mathbf{x}$ contains signals (information) from the environment, which are arranged in order of increasing importance. $\alpha(\mathbf{x})$ can be thought of as a weighted sum of the environmental signals. Each period, the firm

[6] The value of α is normalized to be between 0 and 1 by dividing by $1/2^{N.}$

views an environmental vector $\mathbf{x}$ and uses this information to estimate the value of $\alpha(\mathbf{x})$.

4.2.2 Environmental change

Each period an environmental vector is randomly chosen with probability $1/2^N$. For example, for $N = 10$, each vector has a probability of 0.000977 of being selected. In our simulations it is very unlikely for an environmental vector to be viewed by the firm more than once; the learning process highlights the pattern recognition features of our neural network model of the firm.

4.2.3 Complexity

To measure the complexity of the information processing problem, we define environmental complexity as the number of bits in the vector, N, which, in the simulations below, ranges from a minimum of 5 bits to a maximum of 60.

4.3 The steps

Here we outline the behavior of the firm each period (time subscripts dropped for notational convenience):

1. Each firm observes the environmental state vector $\mathbf{x}$.

2. Based on that each firm estimates a value of the intercept parameter, $\hat{\alpha}_j$. 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. Based on the values the firm estimated in step 2, it makes an output decision using the TFT-type rule given by equation (1).

4. It then observes the true value of α and q_{-j}, and uses this information

to determine its errors using the following rules:[7]

$$\varepsilon_{1j} = (\hat{\alpha}_j - \alpha)^2 \tag{2}$$

$$\varepsilon_{2j} = \left\{ \begin{array}{ll} \left(\hat{q}^j_{-j} - q_{-j}\right)^2 & \text{if } \left(q_{-j} - \frac{\alpha_j}{4}\right) > \rho_j \\ \left(\hat{q}^j_{-j} - \alpha/4\right)^2 & \text{otherwise} \end{array} \right\} \tag{3}$$

Note the contents of equation (3). After the firm observes its rival's output it must select a particular error. If it sees that its rival did not cooperate it therefore yields one error (i.e., the 'did-not-cooperate' one), but if it sees that its rival did cooperate then it uses the other error. This way the firm uses its rival's past performance as guide for the firm's learning.

5. Based on these errors, the firm updates the weight values in its network for improved performance in the next period. This process is outlined in the next section.

5 The firm as artificial neural network

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.[8] The 'nodes' in the managerial and decision layers represent the information processing behavior of agents. Each agent takes a weighted sum of the information it views and applies a type of squashing function to produces an output/signal: a value between zero and one. In the managerial layer, this output is then passed/communicated to the decision layer. Each output node takes a weighted sum of the signals from the managerial layer to produce a particular decision: an estimated intercept value or an estimated rival's output. A graph of the network is shown in Figure 1. This type of ANN is referred to as a backward propagation network (BPN).

[7] In Barr and Saraceno (forthcoming) we demonstrate the how profit and error are negatively related, and that maximizing profit is the same as minimizing the error. For the sake of brevity, we do not show a similar result in this paper since the error rule is relatively more complex, but the same result applies in this case and is confirmed by regression analysis (results not shown).

[8] See Barr and Saraceon (2002; forthcoming) for more details about using ANNs as models of the firm.

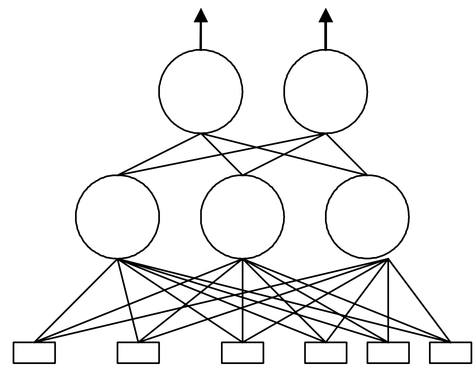

Figure 1: A graph of a neural network.

More specifically, the environmental data (information) layer is a binary vector $\mathbf{x} \in \mathbf{X}$ of length N. Each manager (node) in the management (hidden) layer takes a weighted sum of the data from the data layer. That is, the i^{th} agent in the management layer of firm j calculates

$$in_{ij}^h = \mathbf{w}_{ij}^h \mathbf{x} \equiv \left[w_{1i}^h x_1 + \ldots + w_{Ni}^h x_N \right], \; i = 1, \ldots, M_j; \; j = 1, 2.$$

Thus the set of 'inputs' to the M_j managers of the management layer is

$$\mathbf{in}_j^h = \left(in_{1j}^h, \ldots, in_{ij}^h, \ldots, in_{M_j}^h \right) = \left(\mathbf{w}_{1j}^h \mathbf{x}, ..., \mathbf{w}_{ij}^h \mathbf{x}, ..., \mathbf{w}_{M_j j}^h \mathbf{x} \right).$$

Each manager then transforms the inputs via a squashing (voting) function to produce an output, out_{ij}^h. Here we use one of the most common squashing functions, the sigmoid function: $out_{ij}^h = g(in_{ij}^h) = 1 / \left(1 + \exp\left(-in_{ij}^h\right)\right)$. Large negative values are squashed to zero, high positive ones are squashed to one, and values close to zero are 'squashed' to values close to 0.5.[9]

The vector of processed outputs from the management layer is

$$\mathbf{out}_j^h = \left(out_{1j}^h, \ldots, out_{ij}^h, \ldots, out_{M_j j}^h \right) = \left(g(in_{1j}^h), ..., g(in_{ij}^h), ..., g(in_{M_j j}^h) \right).$$

[9]The sigmoid function can represent the votes of managers since the sigmoid is a continous version of the Heaviside function.

The inputs to the output layer are weighted sums of all the outputs from the hidden layer:

$$in^o_{\iota j} = \mathbf{w}^o_{\iota j}\mathbf{out}^h_j \equiv \left(w^o_{1\iota j}out^h_{1j} + ... + w^o_{M_j\,\iota\,j}out^h_{M_j\,j}\right),\ \iota = 1,2.$$

All weights in both layers can take on any real value. Finally, the outputs of the network–the estimate of demand intercept $\hat{\alpha}_j$ and $\hat{q}^j_{-j}$–are determined by transforming $in^o_{\iota j}$ via the sigmoid function, $\left\{\hat{\alpha}_j = g\left(in^o_{1j}\right), \hat{q}^j_{-j} = g\left(in^o_{2j}\right)\right\}$. We can summarize the behavior of network with two outputs as

$$\begin{aligned} \hat{\alpha}_j &= g\left[\sum_{i=1}^{M_j} w^o_{i1j}\, g\left(\mathbf{w}^h_{ij}\mathbf{x}\right)\right], \\ \hat{q}^j_{-j} &= g\left[\sum_{i=1}^{M_j} w^o_{i2j}\, g\left(\mathbf{w}^h_{ij}\mathbf{x}\right)\right]. \end{aligned}$$

5.1 The learning algorithm

The above process describes the input-output nature of the neural network. However, the distinguishing feature of the network is its ability to learn. After the market price is set and the market clears, the firms learn the true values of the intercept and rival's output. They then use this information to help improve their performance in successive periods. They calculate the error they made and then update the network weights using a gradient decent method that changes the weights in the opposite direction of the gradient of the error with respect to the weight values. We begin with a completely untrained network by selecting random weight values (i.e., we assume the network begins with no prior knowledge of the environment).

An environmental state vector is realized and the networks processes it, as described above, to obtain outputs $\left\{\hat{\alpha}_j,\ \hat{q}^j_{-j}\right\}$. These outputs are compared to true values to get an error for each one, according to equations (2) and (3). Total error is then calculated:

$$\xi_j = \varepsilon_{1j} + \varepsilon_{2j}$$

This information is then propagated backwards as the weights are adjusted according to the learning algorithm, that aims at minimizing the total error, ξ_j. Define $\hat{y}_j = \left\{\hat{\alpha}_j,\ \hat{q}^j_{-j}\right\}$ and $y_j = \{\alpha,\ q_{-j}\}$. The gradient of ξ_j with respect to the output-layer weights is

$$\frac{\partial \xi_j}{\partial w^o_{i\iota j}} = -2\left(y_{\iota j} - g\left(in^o_{\iota j}\right)\right) g'\left(in^o_{\iota j}\right) out^h_{ij},$$

where $i = 1, ..., M_j;\ \iota = 1, 2;\ j = 1, 2$. For the sigmoid function, $g'\left(in^o_{\iota j}\right) = \hat{y}_{\iota j}(1 - \hat{y}_{\iota j})$.

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

$$\frac{\partial \xi_j}{\partial w^h_{ikj}} = -2g'(in^h_{ikj})x_k\left[\left(y_{\iota j} - g\left(in^o_{\iota j}\right)\right) g'\left(in^o_{\iota j}\right) w^o_{i\iota j}\right],$$

where $i = 1, ..., M_j;\ \iota = 1, 2;\ j = 1, 2;\ k = 1, ..., N$.

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. If we define $\delta^o_{\iota j} = .5(y_{\iota j} - \hat{y}_{\iota j})g'(in^o_{\iota j})$, we have the weight adjustment for the output layer as

$$w^o_{i\iota j}(t+1) = w^o_{i\iota j}(t) + \eta\delta^o_{\iota j} out^h_{ij}.$$

Similarly, for the hidden layer,

$$w^h_{ijk}(t+1) = w^h_{ijk}(t) + \eta\delta^h_{i\iota j} x_k,$$

where $\delta^h_{i\iota j} = g'(in^h_{ij})\delta^o_{\iota j} w^o_{i\iota j}$. When the updating of weights is finished, the firm views the next input pattern and repeats the weight-update process.

6 Organization, learning and cooperation: a simulation experiment

In this section we present the results of a simulation experiment. The steps of the experiment were outlined in sections 4 and 5 above. We are mainly interested in two different issues: The first is whether the two firms learn, i.e., if in the long run they are able to map signals from the environment to demand conditions and the opponent's quantity decisions. The second is how the environmental complexity that firms face, especially in the first stages of their interaction, affects their decision to cooperate or defect as well as their profitability. For each set of parameter values, we rerun the simulation 50 times and take average values in order to smooth out fluctuations.[10]

In this section we show some particular runs of the model that are representative of its features. The robustness of these results will then be verified in section 7 by means of an econometric investigation over the parameter space. Here we fix many of the variables and look only at one particular outcome.

6.1 Error

The first question is: Can firms learn both the relationship between the environment and demand and learn to forecast the rival's output decision, given the complexity level of the environment? Below, in Figure 2, we show the results of one firm with $M_1 = M_2 = 8$ nodes in the hidden layer for each firm, $T = 250$ iterations, $N = 10$ inputs, and $\rho_j = 0.05$ $(j = 1, 2)$.[11] We show the separate errors to see that both converge to zero over the long run. We can see that the network is able to improve its forecasts over time. Interestingly, the estimation of the rival's output has a lower error than the estimation of the intercept values.

6.2 Profit and cooperation

Once we made sure the networks learn, we next see how cooperation and profit evolve. In Figure 3 we show how average cooperation rates for firm 1 change over time for two different 'niceness' values. Here we define cooperations rates as follows. Let

$$c_{1t} = \left\{ \begin{array}{c} 1 \text{ if firm 1 plays } C \text{ in period } t \\ 0 \text{ otherwise.} \end{array} \right\}.$$

Then over 50 runs we take the average cooperation rate for each period as $\bar{c}_{1t} = \frac{1}{50} \sum_{r=1}^{50} c_{1rt}$ and this is what we plot versus time in Figure 3. Notice that if $\rho_j = 0$, the cooperation rates steadily decrease over time. This is due

[10]Since firms are adaptive we do not include a discount parameter in the simulations.

[11]For the remainder of this section, unless otherwise noted, we have $T = 250$, $M_1 = M_2 = 8$, $N = 10$.

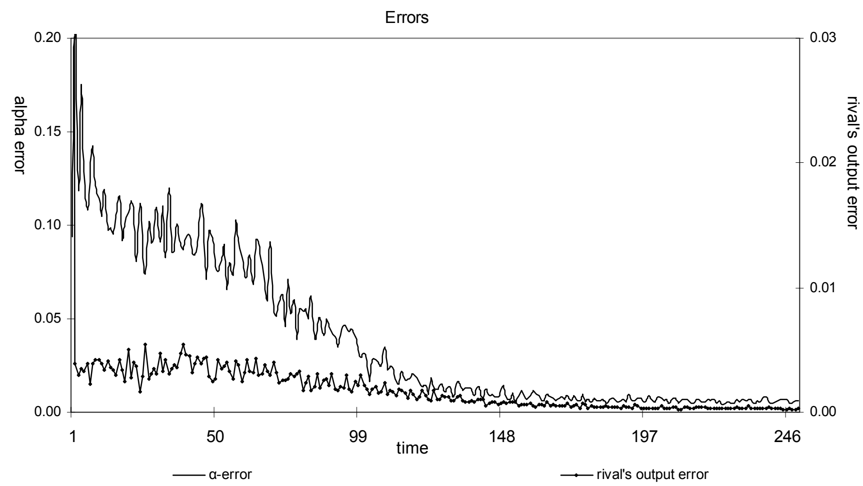

Figure 2: Firm 1 average errors of network over time.

to the fact that the firm continues to make errors in estimating its rival's output and thus is more likely to defect over time. While if the firm has a sufficiently high niceness parameter this helps sustain cooperation in the face of error.

Given how cooperation diverges based on niceness, we can see how profits increase over time; this is shown in Figure 4. We show profits for two different 'niceness values,' $\rho_j = 0$ and $\rho_j = 0.05$ $(j = 1, 2)$ (i.e., both firms play the game with the same niceness parameter). We present the results only for firm 1 since the two firms are symmetric. In the case of $\rho_j = 0$ firms are less willing to cooperate, all else equal. This results in lower profits for the firm. When $\rho_j = 0.05$ profits are relatively higher. Notice, however, that the relative increase in profits due to increased cooperation is not that large. This is due to the fact that since the error is much greater for α, added cooperation only increases profits by a little bit.

A note about 'niceness' We can think of niceness as a measure of 'corporate culture' in the sense that some firms are inherently more aggressive, while some are not. In the population of firms, we are likely to see a wide variation in this inherent firm behavior. In other words, we can think of the degree of niceness as being part of the firm's 'genetic code,' and here

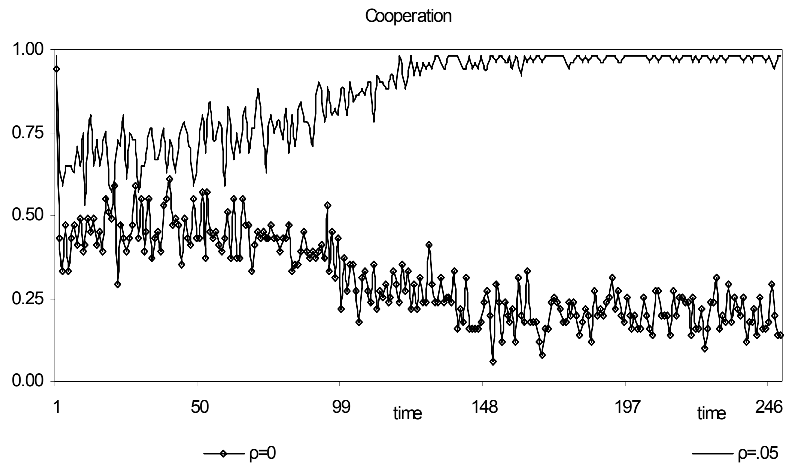

Figure 3: Firm 1 cooperation rates for different 'niceness' values.

we study how differences in firms' niceness values affect the outcome of the game.

6.3 Profits and cooperation versus complexity

Here we look at average profits and average cooperation rates as a function of complexity, given the degree of niceness ($\rho_1 = \rho_2 = 0.05$). We see from Figure 5 that the more information to process (i.e., the greater environmental complexity) the lower is profits and average cooperation. This finding implies that all else equal, the prospects for sustained cooperation is diminished in complex environments.

6.4 Cooperation versus 'niceness'

The above results are given for symmetric firms. Here we ask the question: What does changing the 'niceness' value for firm 1 do to cooperation rates and profits, holding firm two's niceness parameter constant? Figures 6 and 7 show the results. Here we see that if firm 1 increases its niceness but firm two holds its value constant (where $\rho_2 = 0$) then increased niceness does not pay. Notice, in Figure 6, that from a value of about $\rho = 0.25$, firm 1 cooperates almost all the time. The threshold is large enough that defection

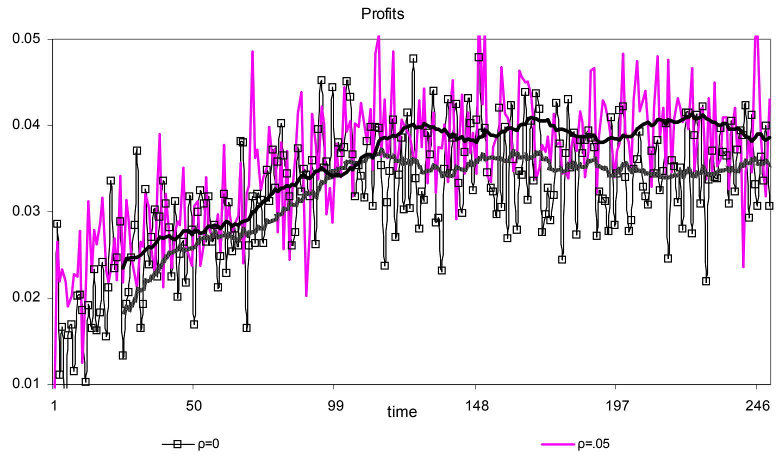

Figure 4: Firm 1 profits over time for two 'niceness' values. Moving Averages (25) also provided.

by the opponent never triggers a non cooperative response. This is why, in section 7, we will focus on a subsample, with $\rho \in [0, 0.25]$.

While firm two also increases its cooperation rate it does not do so at the same rate as firm one. Interestingly firm one experiences a negligible change in profits, while firm two's profits increase due to its 'defection advantage' compared to firm 1.

In summary, the results of this section show that neural networks competing in a duopoly setting can learn both the economic environment and the rival's output decisions. This learning results in increased profit over time, but the level of profit that the two firms can achieve is a function of their willingness to cooperate. The 'nicer' they are the more they can achieve mutual gains to cooperation. But if their niceness threshold is relatively low, they are less likely to cooperate. This is due to the fact that since the environment state changes each period, firms will have some error in estimating the demand function, thus they will always have difficulty estimating the rival's true output and this will make firms more likely to defect. Finally, increasing environmental complexity is associated with lower profits and lower cooperation rates.

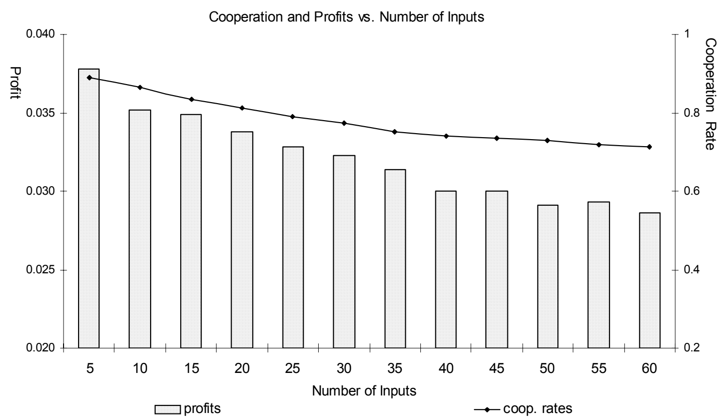

Figure 5: Firm 1 average profits and cooperation rates versus environmental complexity.

7 Regression results

In our model, profits and cooperation depend on a number of variables (environmental complexity, firm complexity, niceness, interaction between firms, etc.) that operate at different levels. This is why, after using particular runs to show some interesting results, we need to check for the robustness of our findings. In this section, we apply standard econometric techniques to a data set generated by random draws of the most interesting parameters. Each observation (we had $3,870$ of them) consists of the parameters and of the two dependent variables: average profit over the run, and the degree of cooperation:[12]

$$\Pi_j = \frac{1}{T}\sum_{t=1}^{T}\pi_{jt};\ C_j = \frac{1}{T}\sum_{t=1}^{T}c_{jt}.$$

The regression results for profit are reported in Table 1. We can give a summary of the results in regards to profits:

- The regression shows a humped-shaped relationship between firm dimension and profits, cet. par. This result is robust as we've also found

[12] The parameters, and the corresponding ranges, are ($j = 1, 2$): $M_j \in [2, 25]$, $N \in [5, 60]$, $T \in [50, 500]$, $\rho_j \in [0, 0.25]$. Notice that, in reporting the results, we only focus on firm one. Firm two is symmetric, so that analyzing it would yield the same results.

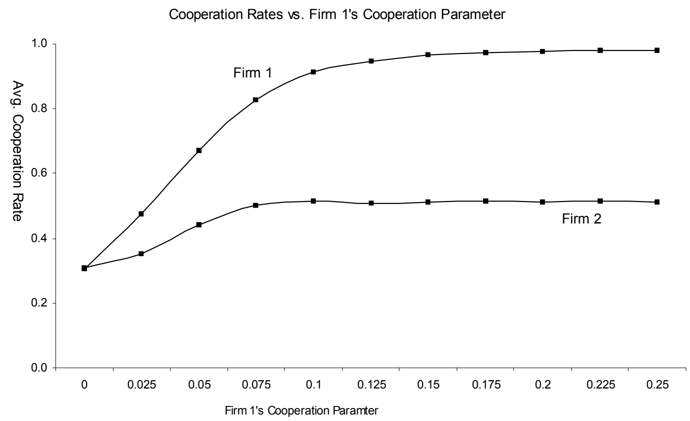

Figure 6: Average cooperation rates for firm 1 and firm 2 versus firm 1's niceness value.

it in our other papers under different circumstances.

- The relationship between profit and environmental complexity is negative and nonlinear.
- Cet. par. increasing time increases profits since it lowers the firm's error.
- Figure 8, where we plot profit as a function of ρ_1 and ρ_2, gives an interesting example of the insights that are possible with our type of analysis. The relationship between profit and own willingness to cooperate is broadly negative, even if non-linearly. Holding constant firm 2's willingness to cooperate, if we increase ρ_1, firm 1 puts itself in the classic prisoner's dilemma situation: if the rival defects then cooperation will yield a lower pay-off. The non-linearity, on the other hand, is linked to uncertainty and learning: extreme types ('harsh' or 'soft') have a more predictable behavior; the opponent will thus be able to learn it better, and act consequently. This will result in higher profits for both firms. Being in the middle, on the other hand, sends contradictory signals to the competitor, and does not favor learning. Such a result would disappear were the firm acting in a world without uncertainty. On the other hand, increased willingness to cooperate from the opponent (larger ρ_2) yields larger profits for firm 1, ceteris paribus.

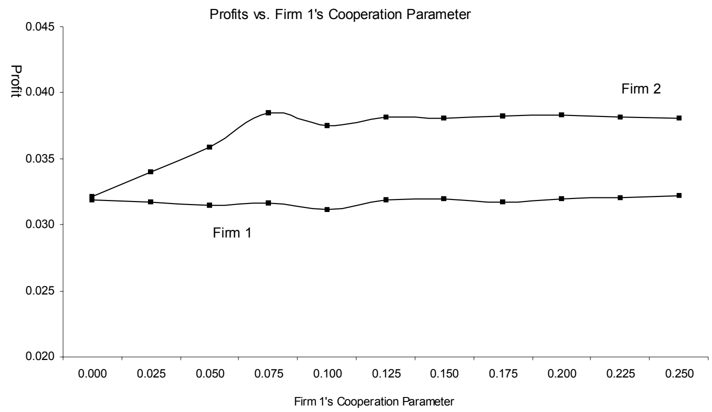

Figure 7: Average profits for firm 1 and firm 2 versus firm 1's niceness value.

Before proceeding, we can draw a first conclusion: Our model confirms previous findings (Barr and Saraceno, 2002; forthcoming) on the relationship between firm profitability, size and the environment it faces; specifically, environmental complexity negatively affects profits, and the trade-off relative to firm size, between speed and accuracy emerges in this setting as well, giving a hump shape relationship between firm size and profit. We also showed that more cooperative firms have higher profits.

We now turn to the other dependent variable, the one specific to this paper: the degree of cooperation; the results of the regression are reported in Table 2.

The regressions yield the following results:

- As expected, the willingness to cooperate by any of the two firms increases the cooperation rate over the relevant range of ρ_1 and ρ_2.
- Also expected is the negative relationship between environmental complexity (N) and degree of cooperation. In a more complex environment, uncertainty makes it harder for a firm to detect defection, by the opponent, who has thus less incentive to cooperate.
- A more interesting picture relates the cooperation rate with firm dimension (see fig 9). We see that after some point, an increased firm size for the rival will induce lower cooperation for firm 1, while we see

Dependent Variable: $10000 \cdot \Pi_1$

Variable	Coeff.	Variable	Coeff.	Variable	Coeff.
M_1	5.94 (0.738)	T^3	0.000 (0.000)	ρ_1	−345.1 (21.8)
M_2	−0.391 (0.108)	T^4	0.000 (0.000)	ρ_2	741.7 (46.1)
M_1^2	−0.825 (0.102)	N	−4.18 (0.327)	ρ_1^2	2,614.5 (200.3)
M_2^2	0.017 (0.004)	N^2	0.172 (0.018)	ρ_2^2	−6,712.7 (736.5)
M_1^3	0.034 (0.006)	N^3	−0.003 (0.000)	ρ_1^3	−5,313.7 (531.1)
M_1^4	−0.001 (0.000)	N^4	0.000 (0.000)	ρ_2^3	26,617 (4399.5)
$M_1 \cdot M_2$	−0.029 (0.003)	$M_1 \cdot T$	0.006 (0.000)	ρ_2^4	−38,600 (8712.4)
T	0.778 (0.049)	$T \cdot N$	−0.003 (0.000)	$\rho_1 \cdot \rho_2$	−134.3 (28.7)
T^2	−0.002 (0.000)	$M_1 \cdot N$	0.053 (0.001)	*cons*	243.9 (3.8)

$R^2 = 0.9351$
Number of obs. 3870

Table 1: Regression results. Dependent variable is profit of firm 1 (standard errors in parentheses; non significant variables have been omitted)

a weaker, but opposite effect for firm 1's size in its own cooperation rates. In general there is a non-linear relationship between own size and cooperation.

To summarize, our regressions substantially confirm the robustness of the results of section 6: cooperation is hampered by more complex environments, and of course by lower "niceness," or willingness to cooperate. In addition, our regression analysis sheds light on the non-linear relationship between firms sizes and cooperation.

8 Network Equilibria

Here we explore the concept of equilibrium with regards to the two important factors that affect performance: the choice of network size and the choice of a niceness level. Above we explored the relationship between performance,

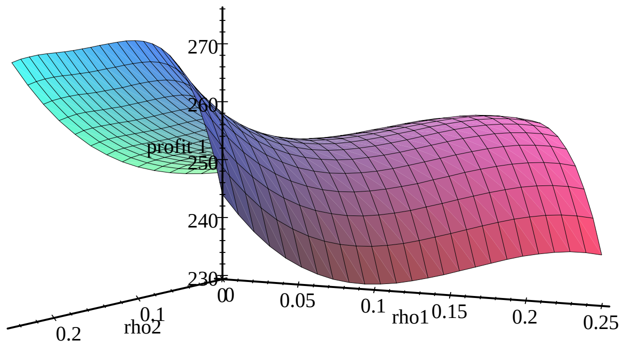

Figure 8: Coefficients taken from the regression presented in table 1.

and cooperation, controlling for several variables including size and niceness. Here we explore in a bit more detail the components of strategic interaction in the sense that a rival's choice of network size and niceness affects the firm's pay-off, as well as cooperation rates. We explore this concept here by looking at the equilibria that emerge over the possible range of network sizes and niceness values.

We define an network equilibrium (NE) as a choice of M_j and ρ_j for each firm such that neither firm has an incentive to change its network size or niceness given its rival's choice of network size and niceness. That is to say, in an equilibrium, each network, given the number of agents (nodes) of its rival and its rival's niceness, finds that switching to another number of agents and/or niceness will decrease its average (total) profit. The equilibrium is a quadruple $\{M_1^*, \rho_1^*, M_2^*, \rho_2^*\}$ such that

$$\Pi_j \left(M_j^*, M_{-j}^*, \rho_j^*, \rho_{-j}^* \right) \geq \Pi_j \left(M_j, M_{-j}^*, \rho_j, \rho_{-j}^* \right), \ \forall M_j, \rho_j \quad j = 1, 2.$$

where

$$\Pi_j = \frac{1}{T} \sum_{t=1}^{T} \pi_{jt} \left(M_j, M_{-j}, \rho_j, \rho_{-j} \right).$$

In this section we ask: What is the relationship between environmental complexity, network size, niceness, profit and cooperation rates at equilibrium? More generally, we are interested in exploring how a firm adapts to

Dependent Variable: $10000 \cdot C_1$

Variable	Coeff.	Variable	Coeff.	Variable	Coeff.
M_1	-22.12 (9.88)	N	-6.27 (1.60)	ρ_2^2	$-92,069.4$ (24232.9)
M_2	105.74 (23.75)	M_1^3	$-.042$ (.020)	ρ_1^3	$1,342,292$ (139338)
M_1^2	1.70 (.81)	N^2	$.049$ ().022	ρ_2^3	$341,911.8$ (144654.6)
M_2^2	-12.92 (3.27)	$M_2 \cdot N$	$.149$ (.040)	ρ_2^4	$-508,592$ (286234)
M_2^3	$.635$ (.176)	$T \cdot N$	$-.019$ (.002)	ρ_1^4	$-1,559,312$ (275964)
M_2^4	$-.0112$ (.003)	ρ_1	$75,181.7$ (1450.4)	$\rho_1 \cdot M_1$	16.56 (10.22)
T	2.92 (.195)	ρ_2	$11,246.4$ (1516.3)	*cons.*	4.043 (83.3)
T^2	$-.002$ (.000)	ρ_1^2	$-463,794.8$ (23296.7)		

$R^2 = .9463$
Number of obs: 3870

Table 2: Regression results. Dependent variable is the average cooperation rate of firm 1 (standard errors in parentheses; non significant variables have been omitted).

given different environmental conditions. That is, for a given environment, the firm seeks to find an optimal network configuration (size and niceness), and this search results in an accompanying profit and cooperation rate.

We focus on the equilibria that exist after T periods, and we do not examine firms changing the number of managers during the learning process. Rather we conduct a kind of comparative statics exercise, whereby we look at the NE that arise for given environmental conditions.

In this experiment, we have networks of different sizes compete for $T = 250$ iterations (for each set of $M's$ and $\rho's$ and input numbers we generate 30 runs and take averages to smooth out fluctuations in each run). That is to say, firms compete for each size, $M_1, M_2 \in \{2, ..., 20\}$, and niceness value $\rho_j = \in \{0, ..., 0.25\}$. We repeat this competition for different levels of complexity, i.e., $N = 5, 10, \ldots, 30$.

For each complexity value we calculated the NEs that emerge (for each number of inputs there were always more than one equilibrium) and then

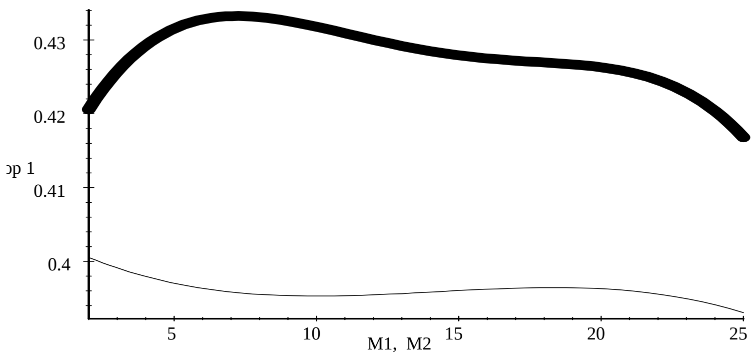

Figure 9: Cooperation rate of firm 1 as a function of M_1 (thin line) and M_2 (thick line). Plot coefficients from table 3.

added the total number of managers of the two firms for each equilibrium to obtain an 'equilibrium industry size' and averaged the $\rho's$ for an average industry niceness. This gave us a data set of 96 NEs (for an average of 16 per complexity level).

Profits and Network Size The first result of this large run is consistent with the findings of the previous sections. Figure 10 shows that in more complex environments the equilibrium structure of the market entails higher firm dimension and lower average profit.

On the other hand, a clear pattern between complexity and cooperation does not appear. The equilibrium average ρ shows a slight increase with complexity, going from 0.078 to 0.086. Interestingly, the degree of cooperation does not seem to be related to the number of inputs (Figure 11).

However, if we delve more deeply into the $\rho's$ and cooperation rates for the two firms, we do see the emergence of some patterns. In equilibrium a polarized behavior structure appears. In general, if one firm evolves towards an aggressive stance (i.e., it selects a low ρ), this induces the competitor to act 'soft,' i.e., to be more accommodating. Figure 12 shows this result.

To investigate more in depth this relationship, we ran a regression with ρ_1 as the dependent variable, and the competitor's niceness level, size, and complexity as independent variables. Table 4 reports the results.

First, notice that the regression confirms the insignificance of environmen-

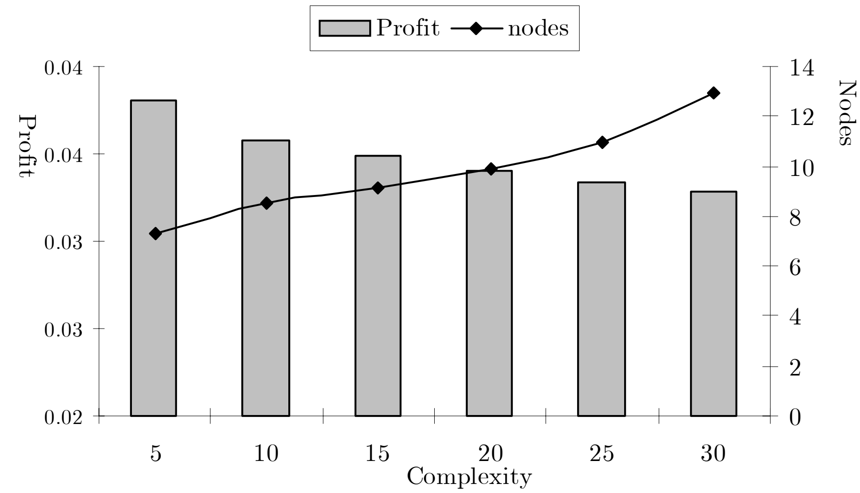

Figure 10: Optimal profit and firm dimension. 'Industry' averages.

tal complexity in explaining cooperative behavior. More surprising, firms' size also appears to be unrelated with equilibrium niceness. Coming to the driving variable, ρ_2, the regression confirms the negative relationship for both small and large values of ρ_2. If we take the coefficients for ρ_2 and plot them against ρ_1 (Figure13), we notice that for values of ρ_2 below about 0.06 the relationship is strongly negative (which means that an aggressive behavior of firm 2 induces soft behavior of 1). Then there is a portion in which the two are complementary, that goes to about $\rho_2 = 0.16$, and then the two become basically unrelated. In other words, once firm two is sufficiently soft, firm 1 becomes aggressive regardless of the value of 2

In summary, we can draw the following conclusions about the network equilibria that emerge:

- Profits in equilibrium are decreasing in complexity.
- Firm size in equilibrium is increasing.
- On, average, cooperation rates do not appear to be affected by complexity; despite the slight positive increase in niceness as we increase complexity. But complexity is strongly related to the niceness values of the two firms.

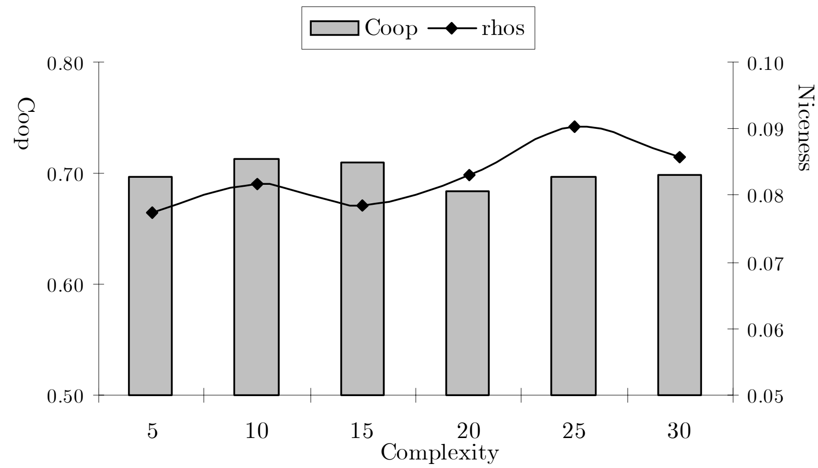

Figure 11: Equilibrium cooperation rates and niceness values. 'Industry' averages.

- Further, firms 'take sides' in their niceness values in equilibrium. In other words, a high nice value by one firm will have a low nice value by another firm in equilibrium.

9 Conclusion

This paper has presented a model of firm learning and cooperation. We investigate the prospects for cooperation given an agent-based model of the firm, which must learn to map environmental signals to changing demand and its rival's output decision. We present two sets of results: (1) outcomes of firm performance and cooperation, holding rival's behavior constant, and (2) the equilibria that emerge due to strategic interaction. In regards to (1), we demonstrate that increased environmental complexity is associated with lower profitability and lower willingness to cooperate. In complex environments frequent cooperation is more sustainable when firms are more willing to be 'nice' in the sense that firms are less likely to defect if they estimate their rival will defect. Further we show that firm size has an effect on both profits and cooperation. Increasing firm size, holding the rival's firm constant, increases profits up to a point and then profits decrease. This trade-off is

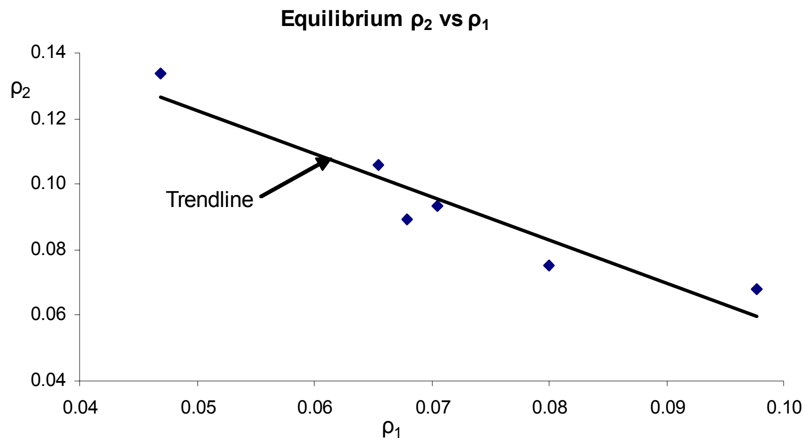

Figure 12: Scatterplot of equilbrium $\rho' s$ (average values for each compleixty level).

due to the fact that adding agents improves accuracy greater than the loss of speed in learning, up to a point. Increasing firm size, ceteris paribus, has an interesting effect on cooperation: we see a non-linear relationship in the firm's cooperation as it increases its size; yet we see that rival's size has a humped-share effect on the firm's cooperation rate.

In regards to network equilibria, we have shown that increasing complexity is associated with larger average firm size and lower profits in equilibrium. We do not find any relationship between cooperation and complexity, but this is due to the fact that firm's 'take sides' in their niceness values. That is to say, in equilibrium, when one firm chooses to be aggressive, the other responds by being relatively nice.

This paper leads to several possible research extensions. First we can explore the prospects for cooperation given that firms play different strategies rather than just the Tit-For-Tat type employed here. Also we can explore the evolution of cooperation given that firms can switch strategies over time. Further we can see how the structure of the network itself affects learning and cooperation. For example, we can study how the networks behave when they have more than one hidden layer or when information to agents is restricted, in order to model specialization.

Dependent Variable: ρ_1

Variable	Coeff.	Variable	Coeff.
ρ_2	-6.47 $(-6.52)^*$	N	-0.0006 (-0.683)
ρ_2^2	76.88 $(4.84)^*$	M_1	-0.0014 (-0.718)
ρ_2^3	-355.9 $(-4.07)^*$	M_2	0.0009 (0.448)
ρ_2^4	567.6 $(3.63)^*$	$const$	0.177 $(5.873)^*$
$R^2 = 0.649$			
Number of obs: 98			

Table 3: Regression results. The dependent variable is the equilibrium niceness value of firm 1. $t - stats$ in parentheses. * Statiscially significant at the 99% level or greater.

References

Axelrod, R. (1984). The Evolution of Cooperation. Basic Books: New York.

Axelrod, R. (1997). The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration, Princeton University Press: Princeton.

Barr, J. and F. Saraceno (forthcoming). Cournot competition, organization and learning. Journal of Economic Dynamics and Control.

Barr, J. and F. Saraceno (2002). A computational theory of the firm. Journal of Economic Behavior and Organization, 49, 345-361.

Carley, K. M. (1996). A Comparison of artificial and human organizations. Journal of Economic Behavior and Organization, 31, 175-191.

Cho, I.-K. (1994). Bounded rationality, neural networks and folk theorem in repeated games with discounting. Economic Theory, 4, 935-957.

Croall, I. F. and J. P. Mason (1992). Industrial Applications of Neural Networks: Project ANNIE Handbook, Springer-Verlag: New York.

Cyert, R.M. and M.H. DeGroot (1973). An analysis of cooperation and learning in a duopoly context. American Economic Review, 63, 24-37.

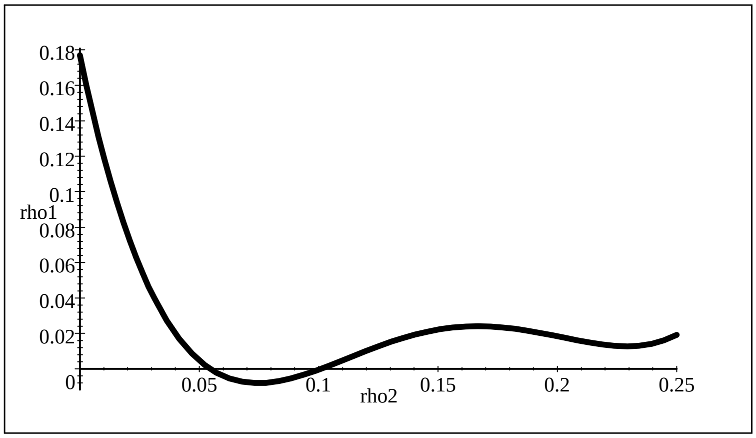

Figure 13: $\rho_1 = 0.177 - 6.47\rho_2 + 76.88\rho_2^2 - 356\rho_2^3 + 567.6\rho_2^4$

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 J. Tirole (1991). *Game Theory*, The MIT Press: Cambridge.

Green, Edward and Robert Porter (1984). Noncooperative Collusion under Imperfect Price Competition. *Econometrica*, 52:87-100.

Ho, T-H. (1996). Finite automata play repeated prisoner's dilemma with information processing costs. *Journal of Economic Dynamics and Control*, 20, 173-207.

Holland, J. H. (1975). *Adaptation in Natural and Artificial Systems*, University of Michigan Press: Ann Arbor.

Kuan C.-M. and H. White (1994). Artificial neural networks: an econometric perspective. *Econometric Reviews*, 13, 1-91.

Miller, J. H. (1996). The coevolution of automata in the repeated prisoner's dilemma. *Journal of Economic Behavior and Organization*, 29, 87-112.

Radner, R. (1993). The organization of decentralized information processing. *Econometrica*, 61, 1109-1146.

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

Skapura, D. M. (1996). *Building Neural Networks*, Addison-Wesley: New York.

Tirole, J. (1988). *The Theory of Industrial Organization,* The MIT Press: Cambridge.

Verboven, F. (1997). Collusive behavior with heterogeneous firms. *Journal of Economic Behavior and Organization,* 33, 122-139.

Vriend, N. J. (2000). An illustration of the essential difference between individual and social learning, and its consequences for computational analyses. *Journal of Economic Dynamics and Control*, 24, 1–19.