

Organization, Learning and Cooperation*

Jason Barr[†] and Francesco Saraceno[‡]

Forthcoming in the *Journal of Economic Behavior and Organization*

Abstract

This paper models the organization of the firm as a type of artificial neural network in a duopoly setting. The firm plays a repeated Prisoner's Dilemma type game, but must also learn to map environmental signals to demand parameters and to its rival's willingness to cooperate. We study the prospects for cooperation given the need for the firm to learn the environment and its rival's output. We show how profit and cooperation rates are affected by the sizes of both firms, their willingness to cooperate, and by environmental complexity. In addition, we investigate equilibrium firm size and cooperation rates.

JEL Classification: C63, C72, D21, D83, L13

Key words: Artificial Neural Networks, Prisoner's Dilemma, Cooperation, Firm Learning

*An earlier version of this paper was 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, Brendan Cunningham and seminar participants at the U.S. Naval Academy for comments and suggestions. The paper also greatly benefited from the comments of two anonymous referees. The usual caveats apply.

[†]Corresponding Author: Department of Economics, Rutgers University, Newark, NJ 07102. email: jmbarr@rutgers.edu, ph: 973-353-5835.

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

1 Introduction

Information processing and decision making by firms are typically not done by one person. Rather decisions are made by groups of people either in committees or hierarchical structures. Bounded rationality and computational costs preclude the possibility of any one agent collecting, processing and deciding about information relevant to the firm and its profitability. 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, 2005), we model the firm as a type of artificial neural network (ANN) that must make output decisions in a Cournot duopoly framework. 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 as non-linear forecasting equations (Kuan and White, 1992).

Here we use ANNs as a model of a firm’s organization, which allows us to make explicit the nature and the costs 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; 2005), the benefit to the firm of increased resources devoted to information processing (IP) is better knowledge of the environment; but the costs include the wage of IP agents, and, especially, the time costs involved with processing and communicating this information.

Barr and Saraceno (2005) investigate how learning affects a firm’s ability to produce along the best response function in a Cournot duopoly game with unknown demand. In that setting, firms choose an output level after observing the state of the environment. After the market clearing price reveals the true position of demand and the opponent’s output choice, firms compare their output choices to the best response they should have played. We show that ANNs can learn to converge to the Nash equilibrium output, given the environmental complexity. Further we show how environmental complexity affects both optimal organizational size and profits. In Barr and Saraceno (2005) we do not consider the possibility of collusion, in the sense that firms could possibly learn to restrict their outputs, thus gaining larger profits.

The difficulty of learning the environment affects a firm’s ability to coop-

erate because in complex environments a firm will have trouble distinguishing whether the variation in the market clearing price is coming from changing environmental characteristics or from the variation in their rival's strategy. This problem has been the subject of several works (see, for example, Green and Porter, 1984). We study a similar problem in the framework of adaptive, multi-agent firms who have to learn how the environment affects their demand curves, and whether the other duopolist is willing to cooperate or not.

While U.S. antitrust law prohibits outright collusion, the temptation to collude remains strong. In fact, price-fixing cases continues to be the most frequent type of case prosecuted by the U.S. Department of Justice (Viscusi, *et al.*, 2001). Under current antitrust law, proving price-fixing requires the presence of a "smoking gun," such as price-fixing agreement documents; tacit collusion cannot be prosecuted under current law (Viscusi, *et al.*, 2001). Thus firms have the incentive to try to raise prices tacitly by sending signals to other firms about their willingness to cooperate to restrict output and raise prices. We are interested in understanding what types of competitive environments can foster tacit collusion.

In our model two neural networks compete in a repeated duopoly setting. Each period, the firms (networks) view an environmental state vector (an N -length vector of 0's and 1's) which they use to estimate two variables: the intercept parameter of the demand curve and their rival's output. Each firm uses its estimate of its rival's output to guess whether it is going to defect or cooperate. This guess is also determined by its own predisposition to cooperate, which we capture by means of a "niceness" parameter. If the estimate of the rival's output is greater than the shared-monopoly output plus some margin depending on its "niceness," 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 will defect if I think you are going to defect, otherwise I will cooperate."

After choosing an output, the firm observes the market clearing price, the rival's output and the true demand intercept, which allows it to compute estimation errors; these are then used to improve its performance in subsequent periods. Thus the knowledge of both the environment and the rival's past decisions lie within the network itself, rather than any one agent; the network serves as a economical knowledge storage device. The questions we ask are:

- 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?
- How are a firm’s profits and cooperation affected by “niceness,” i.e., the firm’s willingness to forgive deviations from the cooperative output, and its rival’s size and niceness?

To anticipate some of our results, first we show that firms in this setting are able to learn, and they converge to the equilibrium of the Cournot game. Next we show that environmental complexity and firm dimension interact in a complex way in regard to both profits and cooperation. Performance and cooperation depend on, not only the difficulty of the IP task, but also firm sizes and the willingness of each firm to cooperate. Further, we show that environmental complexity affects, in equilibrium, total industry size (defined as the number of agents in the two firms), but has mixed effects on cooperation.

In the next section we briefly review the literature to which our work relates. Section 3 outlines the standard Cournot model that we use, which is an extension of the repeated Prisoner’s Dilemma 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 to 8 present the results of our simulation experiments. Finally, in section 9 we conclude.

2 Related Literature

There is a rich literature on the issue of cooperation and defection in Prisoner’s Dilemma games, of which the Cournot game is a variant. In this general framework, two players must make a decision over 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 there can be mutual gain.

In terms of the repeated Prisoner’s Dilemma (RPD), a standard theoretic result is the “folk theorem,” which says that if agents are patient enough, then there is an infinite number of Nash Equilibrium outcomes that have higher payoffs 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 over time if automata machines adapt using a genetic algorithm (Holland, 1975) that allows the strategic environment to change. 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. Our focus is a little different in that we investigate the role that network structure and environmental complexity play in the outcome of the game. For simplicity we limit our attention to a type of tit-for-tat strategy, and focus on the learning process, and on the relationships between firm size/complexity, environmental complexity, profitability and cooperation.

Cournot competition is a common application of RPD games (see Tirole, 1988, for example). In these models collusion is possible if firms are sufficiently patient and the threat of punishment exists (Verboven, 1997). Cyret and DeGroot (1973) show that firms can 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 on whether agents

perform individual rather than social learning. Our model is similar to these papers in the sense that firms are adaptive, but we also look at the effect of environmental complexity and the (heterogeneous) willingness to cooperate (the “niceness” parameter).

This paper is also related to recent work on agent-based economics, in two different senses of the term.¹ First, our conception of the firm as an artificial neural network fits within the set of recent models discussed by Chang and Harrington (2006), who describe agent-based computational models of firms as having the following properties (pp. 1276-77): the firm is modeled as a network of individual agents; these agents typically process information and transmit this information to others within the network; the firms engage in adaptive search methods rather than simply playing an equilibrium response; and the combination of the firm-as-network with adaptive learning prohibits the presentation of closed-form solutions.

Similar work in this vein include Carley (1996), and DeCanio and Watkins (1998). In general, these papers study which types of networks minimize the costs of processing and communicating information. However, our work is different from these models in that we directly model the relationship between the external environmental variables, firm learning and performance. In addition, we explicitly apply this framework to a model of Cournot competition and cooperation, which to our knowledge has not been done before Barr and Saraceno (2005).

Another set of agent-based models include interacting, heterogenous agents, such as those presented in Axelrod (1997). Our firms-as-agents also display a degree of inherent heterogeneity in their “types,” or willingness to cooperate, as well in their size. Similarly to Axelrod (1984), we show that the “personality” of a player’s strategy is an important determinant of the sustainability of cooperation.

Henrich, *et al.* (2001) show that players’ willingness to cooperate and reciprocate varies widely across societies, and is deeply embedded in the cultures of the players. Organizations, too, have different “personalities,” and corporate cultures (Camerer and Vepsäläinen, 1988; Casson, 1991), and firm behavior and decisions vary, even within the same industries where the technological, market and institutional structures are often the same (see for

¹Arguably there is a third definition of agent-based economic models, not discussed here, that tries to “grow” economies from the ground up, where simple rules of behavior can generate complex outcomes. The work of Axtell and Epstein (1996) fits within this category. See Tesfatsion (2002) for review of these models.

example, Hannon and Freeman, 1989 and Nelson and Winter, 1982).

3 The Duopoly Framework

This section gives a brief summary of the standard Cournot game. 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

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

For the moment, we take $\alpha_t = \alpha$ as constant and known to the firms. We also assume that the slope is constant and normalized to one. Profits for each firm are

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

where c_j is costs, such as the cost of network, and without loss of generality it is set to zero for convenience. Under the standard Cournot assumptions, the best response function is given by

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

with a Nash Equilibrium quantity and profit of

$$q^{ne} = \frac{\alpha}{3}, \quad \pi^{ne} = \frac{\alpha^2}{9}.$$

If the two firms could coordinate their output decisions and act as a monopoly their profit from joint production Q would be

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

Assuming the firms share production and profits equally, each would have an output and profit of

$$q_j^m = \frac{\alpha}{4}, \quad \pi_j^m = \frac{\alpha^2}{8} > \pi_j^{ne} = \frac{\alpha^2}{9}.$$

This is a typical prisoner dilemma's game. In a single shot game, the cooperation outcome is not an equilibrium; 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:

$$q_j^d = \frac{1}{2} [\alpha - q^m] = \frac{3}{8}\alpha$$

$$\pi_j^d = \frac{9}{64}\alpha^2 > \pi_j^m = \frac{\alpha^2}{8}.$$

In in a repeated framework with either infinite or undefined time horizon, however, cooperation may emerge and be sustainable. 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, $\Pi_j = \sum_{t=1}^{\infty} \delta^{t-1} \pi_{jt}$.² In this case it can be shown that if the discount rate is sufficiently large, almost any strategy involving cooperation (including the tit-for-tat strategy described in the next section) can be an equilibrium strategy. This is, in fact, what the folk theorem mentioned above says: any feasible expected payoff can be sustained in an equilibrium as long as each player can expect a payoff larger than the noncooperative one. A backward induction argument, on the other hand, can show 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 this paper is the consideration of the effects of uncertainty on the sustainability of cooperative equilibria (as in Green and Porter, 1984). If the intercept term of a linear demand function shifts according to a given probability distribution, then 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 enforce cooperation needs to be more complex. This typically involves a tradeoff 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 partner may be tempted to cheat. To sustain cooperation firms have to punish their opponents only if prices and profits deviate "too much" from the cooperative level. We adopt a similar perspective in what follows, with a crucial difference: uncertainty can be reduced through firm learning. Our focus, in fact, is on this learning process, and on how it affects the prospect for cooperation. Notice, too, that the

²Notice 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.

patience necessary to maintain cooperation in the deterministic model becomes less relevant in the stochastic case, where environmental effects are relatively more important. This is shown by Levenstein and Suslow (2006), who discuss the effect of demand changes have on cartel duration.

4 A Model of Firm Learning in a Repeated Cournot Game

4.1 Firm Strategies

In this paper each network adopts a variant of the tit-for-tat (TFT) strategy. The standard TFT strategy says that a firm should begin by cooperating, and then use the same behavior as the rival's prior action. This strategy, in general, conforms to the four rules-of-thumb discussed by Axelrod (1984) for strategies that are likely to promote cooperation among boundedly rational agents. (1) Be nice: never be the first to defect. (2) Be forgiving: be willing to return to cooperation even if your opponent defects. (3) Be simple: the easier it is for your rival to discover a pattern in your behavior, the easier it is for him to learn to cooperate. (4) Don't be envious: don't ask how well you are doing compared to your rival, but rather try to see how you can improve your own performance.

In industrial organization, collusion has a long history, and forms the basis, for example, of U.S. antitrust law and competition policy. Since outright collusion is illegal, firms may find it in their best interest to signal their willingness to collude. A simple way two model this possibility is to have two firms maintain a simple strategy, such as TFT. The issue then becomes: given that each firm has a willingness to collude, what are the environmental and organizational factors that determine whether collusion is a feasible and sustainable outcome of the game? Here, for example, a grim-trigger or related strategy would simply not work due to the error that a firm makes when evaluating a rival's actions. In addition, a more sophisticated strategy is likely to produce too much "noise" to generate cooperation. In our framework TFT allows for the possibility of cooperation once firms learn to distinguish the part of price variation that is due to environmental change versus the part due to their rival's output decisions. We leave the introduction of evolving strategies for future work.

We employ, however, a slight variation of the TFT strategy. Since firms

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

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

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 a firm estimates its rival to be a cheater, i.e., to deviate from forecast monopoly profit ($\hat{q}_{-j}^j > \hat{\alpha}_j/4 + \rho_j$), then it plays the forecasted Cournot output; that is, it defects as well.

The threshold value $\rho_j \geq 0$ represents, as already mentioned, a firm's niceness. For relatively small values of ρ_j , firm j will defect relatively more often; for large enough values of ρ_j , the firm will be so nice that it will never defect. Notice that in making this decision, as in Green and Porter (1984), the firm has two possible sources of error: the first is the environment, and the second is the opponent's quantity; this uncertainty explains why it will allow a deviation ρ_j from the monopoly output before reverting to the noncooperative quantity. We treat ρ_j as a parameter, and we look at the evolution of cooperation for exogenous degrees of niceness.³ We allow ρ_j to vary across firms, so that it becomes an additional source of heterogeneity.

4.2 The Economic Environment

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

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

where x_k is the k^{th} element of \mathbf{x} (time subscripts are dropped for notational convenience). This functional relationship converts a binary digit vector into its decimal equivalent. The value of α is normalized to be between 0 and 1 by dividing the sum by $2^N - 1$. $\alpha(\mathbf{x})$ can be thought of as a weighted

³ ρ_j could be modeled as endogenous (i.e., the history of the game affects not only a agent's cooperative behavior, but also its general willingness to cooperate). We leave this extension for future work.

sum of the environmental signals, which are arranged in order of increasing importance. Each period, the firm views an environmental vector \mathbf{x} and uses this information to estimate the intercept of the demand curve.

Each period an environmental vector is randomly chosen with probability $1/2^N$.⁴ To measure the complexity of the information processing problem, we define environmental complexity as the number of bits in the vector, N ; the more elements that contribute to determine the position of the demand curve, the more complex is the environment. In the simulations below N ranges from a minimum of 5 bits to a maximum of 50.

4.3 The Sequence

Each period:

- Firms observe an environmental state vector \mathbf{x} .
- Then, based on \mathbf{x} , each firm ($j = 1, 2$) estimates the value of the intercept parameter, $\hat{\alpha}_j$, and their rival's choice of output, \hat{q}_{-j}^j .
- Next, $\hat{\alpha}_j$ and \hat{q}_{-j}^j serve as the basis for each firm's output decision using the TFT-type rule given by equation (1).
- Firms then observe the true value of α and q_{-j} , and use this information to determine their errors, using the following rules:⁵

$$\varepsilon_j^\alpha = (\hat{\alpha}_j - \alpha)^2, \quad (2)$$

$$\varepsilon_j^q = \left\{ \begin{array}{l} (\hat{q}_{-j}^j - q_{-j})^2 \text{ if } (q_{-j} - \frac{\alpha_j}{4}) > \rho_j \\ (\hat{q}_{-j}^j - \frac{\alpha}{4})^2 \text{ otherwise} \end{array} \right\}. \quad (3)$$

Note the contents of equation (3). If the rival's output is close enough to the shared monopoly value, the firm will assume that the rival cooperated and will compute the error as a deviation from $\alpha/4$. Otherwise

⁴For example, for $N = 10$, each vector has a probability of 0.000977 of being selected. In the computer science literature, the ANN usually "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 unseen data. Here, the networks train and generalize simultaneously.

⁵Barr and Saraceno (2005) show 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 the equivalent result in this paper.

it will assume that the rival did not cooperate, and will compute the error accordingly. This way the rival’s past choices are used as a guide for the firm’s learning.

- 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.

Notice that each firm must estimate the intercept of the demand function—which directly depends on the observed environmental vector—and the rival’s output choice—which is related to the environment only indirectly via learning. This feature of the model captures the fact that a firm cannot generally acquire information about its rival’s strategy (short of industrial espionage) until it is actually implemented. In most markets, decisions must be made beforehand, and therefore must rely only on observable environmental information.

Though the environmental state does not directly map to output decisions, the networks can still learn the rival’s behavior from observing the environmental state. This happens because the firms follow the same general rules of behavior (i.e., cooperate if estimated output choice of the rival is not too large, defect otherwise) and thus using the environmental signals is a reliable indicator of what the rival will do, given that each firm “trains” to learn its rival’s behavior over time.⁶

5 The Firm as an Artificial Neural Network

The neural networks that we use as models of the firm are comprised of three “layers”: the input layer (i.e., the environmental state vectors), a hidden/managerial layer, and an output/executive layer. The nodes in the managerial and decision layers represent the information processing behavior of agents. The number of managers, M_j , captures the information processing capacity of the firm (in Computational Learning Theory it is referred to as “hypothesis complexity”; see Niyogi, 1998, p. 23). Hereafter we will call it the size of the network/firm.

⁶In future work we could also envision a case where rivals use different strategies (or even evolving ones), and another agent (output node) can be dedicated to learning, for example, some information about the choice of a rival’s strategy, which could then help a firm in deciding on an output choice.

Each agent takes a weighted sum of the information it views and applies a type of squashing function to produce an output/signal—a value between zero and one. In the managerial layer, this output is then communicated to the executive 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.




Figure 1: Graph of a neural network.

Our choice of ANNs as a models of the firm is motivated by the fact that the behavior of the network has many features in common with the information processing and decision making activities of an organization: (1) The firm is a coordinated network of agents (nodes), which processes information in a decentralized manner, with both parallel and serial computations. (2) Organizations have hierarchical structures, and thus decisions are made within decisions. (3) Firms learn by experience; agents become proficient via trial and error. (4) Organizations must adapt to their environments. That is to say, firm performance is determined by how well firms can “fit” within the world in which they operate. (5) As agents process information, they learn to recognize patterns and, therefore, the firm is able to generalize to new information. (6) The firm is a coordinated collection of specialized agents. Thus the knowledge of the firm is greater than the sum of the knowledge of its agents. (7) Firms that operate in competitive environments generally have limited resources to devote to information processing, and this competition imposes time constraints. In short, the network can be thought of as

a simple, stylized model of the “brains” of a firm’s operations

More accurate learning implies a better understanding of the environment and more profitable decisions in general, but this accuracy requires more agents and time. Thus competitive firms face a tradeoff between speed and accuracy. This tradeoff is also a typical feature of ANNs and of “learning machines” in general (Niyogi, 1998). On one hand a large network can asymptotically achieve a very low error, even if the rate of convergence to the minimum possible error is relatively slow.⁷ On the other hand, a small network will converge much faster to its minimum attainable error, that nevertheless will be larger. For a given level of environmental complexity, it is possible to define an optimal network size that balances speed and accuracy. In a relatively simple environment, as will be demonstrated below, a small network is the one that best balances this tradeoff since speed is relatively more important; while as the environmental complexity grows, a larger network is optimal.

Network j produces its output as follows. Each manager (node) in the management (hidden) layer takes a weighted sum of the data from the environmental layer (of size N). That is, the i^{th} agent in the management layer of firm j calculates

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

Thus the set of “inputs” to the M_j managers of the management layer is

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

Each manager then transforms the inputs via a squashing function to produce an output, out_{ij}^h . Here we use one of the most common squashing functions, the sigmoid: $out_{ij}^h = g(in_{ij}^h) = 1 / (1 + e^{-in_{ij}^h})$. 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.

The vector of processed outputs from the management layer is

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

⁷White (1992), for example, has shown that a neural network with an infinite number of nodes is a universal approximator.

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

$$in_{\iota j}^o = \mathbf{w}_{\iota j}^o \mathbf{out}_j^h \equiv \left(w_{1\iota j}^o out_{1j}^h + \dots + w_{M_j \iota j}^o out_{M_j j}^h \right), \iota = 1, 2.$$

Finally, the outputs of the network— $\hat{\alpha}_j$ and \hat{q}_{-j}^j —are determined by transforming $in_{\iota j}^o$ via the sigmoid function, $\{\hat{\alpha}_j = g(in_{1j}^o), \hat{q}_{-j}^j = g(in_{2j}^o)\}$. We can compactly write the input-output nature of the neural network as

$$\begin{aligned} \hat{\alpha}_j &= g \left(\sum_{i=1}^{M_j} w_{i1j}^o g(\mathbf{w}_{ij}^h \mathbf{x}) \right), \\ \hat{q}_{-j}^j &= g \left(\sum_{i=1}^{M_j} w_{i2j}^o g(\mathbf{w}_{ij}^h \mathbf{x}) \right). \end{aligned}$$

A distinguishing feature of the network is its ability to learn. In this paper we use the most popular type of learning rule, the Backward Propagation Network (BPN) algorithm, which owes its name to the fact that learning takes place by updating the weights at the output layer first and then “backwards” to the input layer.

Here, after the price is determined and the market clears, firms learn the true values of α and q_{-j} . They then use this information to calculate the errors, and update the network weights using a gradient decent algorithm, which 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 in the range $w \in [-0.1, 0.1]$. The range of the initial draws, of course, affects the speed of learning, which is crucial in our setting, but by using the same range for all networks, we make sure that relative learning performances of different networks/firms are unaffected. Using different weight ranges does not affect the qualitative results presented below.

The outputs $\hat{y}_j = \{\hat{\alpha}_j, \hat{q}_{-j}^j\}$ are compared to true values to get an error for each one, according to equations (2) and (3). Total error ξ_j is then calculated as

$$\xi_j = \varepsilon_j^\alpha + \varepsilon_j^q.$$

This information is then propagated backwards as the weights are adjusted according to the learning algorithm, which aims to minimize ξ_j . Define

$y_j = \{\alpha, q_{-j}\}$. The gradient of ξ_j with respect to the output-layer weights is

$$\frac{\partial \xi_j}{\partial w_{ij}^o} = -2 (y_{\iota j} - g(in_{\iota j}^o)) g'(in_{\iota j}^o) out_{ij}^h,$$

where $i = 1, \dots, M_j$; $\iota = 1, 2$; $j = 1, 2$. Similarly, we can find the gradient of the error surface with respect to the hidden layer weights:

$$\frac{\partial \xi_j}{\partial w_{ikj}^h} = -2g'(in_{ikj}^h)x_k [(y_{\iota j} - g(in_{\iota j}^o)) g'(in_{\iota j}^o) w_{ij}^o],$$

where $k = 1, \dots, N$. Once the gradients are calculated, the weights are adjusted a small amount in the opposite (negative) direction of the gradient. The constant η , the learning-rate parameter, is introduced to smooth the updating process.⁸ If we define $\delta_{ij}^o \equiv \frac{1}{2}(y_{\iota j} - \hat{y}_{\iota j})g'(in_{\iota j}^o)$, we have the weight adjustment for the output layer as

$$w_{ij}^o(t+1) = w_{ij}^o(t) + \eta \delta_{ij}^o out_{ij}^h.$$

Similarly, for the hidden layer,

$$w_{ijk}^h(t+1) = w_{ijk}^h(t) + \eta \delta_{ij}^h x_k,$$

where $\delta_{ij}^h \equiv g'(in_{ij}^h)\delta_{ij}^o w_{ij}^o$. When the updating of weights is finished, the firm views the next input pattern and repeats the weight-updating process.

The backward-propagation rule is a reasonable model of organizational learning. The rule says that in each period the agents in the output layer compare their estimates to the actual values; the information from the error is propagated backwards, so that information is transmitted down the hierarchy. We can think of this as if the executives are sending a memorandum to the other agents, telling them how they can help improve the firm's performance over time.

Note that the networks are "learning from scratch," since the weights are initially chosen at random. Thus in the first few periods, output by the network will essentially be guesses; it takes time for the network to learn how recognize the underlying patterns and map the inputs to the outputs with reduced error. Below we look at the average error over time (specifically, over 250 weight-adjustment iterations) because we are interested in how well the firm balances the tradeoff between speed and accuracy.

⁸For the simulations, η is fixed at 10.

6 Organization, Learning and Cooperation: A Simulation Experiment

This section presents the results of a simulation experiment.⁹ 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, and how well, the two firms can 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 environmental complexity, especially in the early stages of firm interaction, and network sizes affect the decision to cooperate or defect and the profitability. For each set of parameter values, we rerun the simulation 50 times and take averages in order to smooth out differences due to the initial random weights.

Here we show some particular runs of the model that are representative of its features. The robustness of these results will be verified in section 7 by means of an econometric investigation over the parameter space.

Learning

The first question is whether firms are able to learn over time the relationship between the environment and both demand and the rival's output decision. In figure 2 we show the errors of firm 1.¹⁰ Both errors are normalized ($\sqrt{\varepsilon_{1,t}^\alpha}/\alpha_t$ and $\sqrt{\varepsilon_{1,t}^q}/q_{2,t}$) to make them comparable. We can see that the network is able to improve both forecasts over time. The figure also shows that, as would be expected, $\sqrt{\varepsilon_{1,t}^\alpha}/\alpha_t$ converges to a lower level than $\sqrt{\varepsilon_{1,t}^q}/q_{2,t}$. This is explained by the fact that the rival's output is only indirectly related to the environment.

Cooperation

Next, we see how cooperation and profits evolve. We define the *cooperation rate* as follows. First, let

$$c_{jt} = \left\{ \begin{array}{l} 1, \text{ if firm } j \text{ plays the shared monopoly in period } t \\ 0, \text{ if firm } j \text{ plays the best response in period } t \end{array} \right\}.$$

⁹The simulations were performed in Mathematica. The code is available upon request.

¹⁰For the remainder of this section, unless otherwise noted, we have $N = 10$, $T = 250$, $M_j = 8$, $\rho_j = 0.05$, $j = 1, 2$.

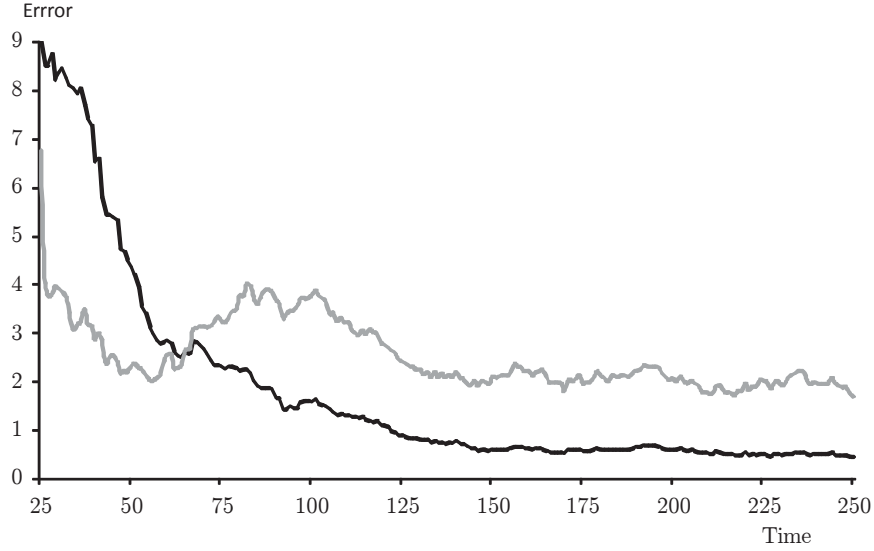


Figure 2: Errors of firm one over time. Moving average (MA(25)) of $\sqrt{\varepsilon_{1,t}^\alpha}/\alpha_t$ (black line) and $\sqrt{\varepsilon_{1,t}^q}/q_{2,t}$ (light line).

Then, over 50 runs we take the cooperation rate for each period as $\bar{c}_{jt} = \frac{1}{50} \sum_{r=1}^{50} c_{1rt}$. We plot \bar{c}_{jt} versus time in figure 3 for two different “niceness” values. If $\rho_1 = \rho_2 = 0$, the cooperation rate steadily decreases over time: errors are systematically attributed to the rival’s unwillingness to cooperate, and the firm will as a consequence defect most of the time. If the firm has a sufficiently high niceness parameter, on the other hand, increasing proficiency in forecasts (learning) will be associated with increasing rates of cooperation. Thus, sustained cooperation in this model becomes possible if firms are sufficiently tolerant facing deviations from the collusive outcome during the learning phase.

Profits and cooperation versus complexity Complexity affects learning and, therefore, cooperation rates and profits. We define average profit and average cooperation over the $T = 250$ iterations as $\Pi_j = T^{-1} \sum_{t=1}^T \pi_{jt}$ and $C_j = T^{-1} \sum_{t=1}^T \bar{c}_{jt}$. Figure 4 plots Π_j and C_j for several complexity levels (N). The figure shows that all else equal, for given network sizes, the prospect for sustained cooperation is diminished in complex environments.




Figure 3: Evolution of firm one’s cooperation rates (MA(25)) for $\rho_1 = \rho_2 = 0$ (black line) and $\rho_1 = \rho_2 = 0.05$ (light line).

This is quite intuitive, as increased complexity makes learning harder, and firms tend to impute undesired outcomes to noncooperative behavior from their rival. As a consequence, they will tend to punish/defect more often.

7 Regression Results

Profits and cooperation choices depend on a number of variables, that interact at different levels (environment, firm complexity, interaction between firms). This is why, after using particular runs to show interesting results of our model, we need to check for the robustness of our findings. This section will apply standard econometric techniques to a data set generated by our model using random draws of the most interesting parameters. Each observation of the data set was built as follows: we randomly and independently selected a set of parameters from the following parameter space: $M_j \in [2, 25]$, $N \in [5, 50]$, $T \in [50, 500]$, $\rho_j \in [0, 0.2]$. Then we let the two firms compete for T iterations, and we recorded the average profit and cooperation rates, Π_j and C_j . Notice that, because of the way we constructed the experiment,




Figure 4: Firm one’s average profits and cooperation rates versus environmental complexity. $M_1 = M_2 = 8$, $\rho_1 = \rho_2 = 0.05$, $N = 10$.

firms are symmetric; we only report the results for firm one.

Given that we are working with artificial data, it is important to verify that the standard ordinary least squares assumptions, and hence the Gauss Markov Theorem, apply to our analysis (see Wooldridge, 2003). The first assumption is that the underlying population regression function is linear in the coefficients. This assumption is valid in our case, since we use high order polynomial terms with linear coefficients to fit the data. As is a standard result, a polynomial of high enough order can approximate any function with any specified degree of accuracy. Next, the method we used to generate the experiment guarantees that the exogenous variables are a randomly determined, unbiased sample. The random drawing of parameters from uniform distributions also guarantees that errors have zero mean and no serial correlation. In addition, we can further assume that errors are asymptotically normal because of the randomness involved with both the initial weight generation and the random drawing of the environmental states.

Finally, we faced a dilemma about the choice of regressors. Given the highly complex nature of the data, we found that using the standard tests for functional form—adjusted R^2 , the Aikake Information Criterion and Ram-

sey’s RESET test—we would not reject a polynomial of relatively high order (for example, the coefficient on M_1^8 was statistically significant at the 99% confidence level). However, much simpler regressions gave very similar results in terms of goodness of fit (the simple regression had an R^2 of 0.94 for Π_1 and of 0.925 for C_1 ; the complex regression had R^2 of 0.96 and 0.94, respectively). Including many higher order polynomials did not materially increase the explanatory power of the regression and potentially over-fit the data. Furthermore, including these additional terms made essentially imperceptible changes in the effect of the exogenous variables on the dependant variables. For these reasons we present here the results of the relatively simple regressions.

Dependent Variable: 10,000 · Π_1					
Variable	Coeff.	Variable	Coeff.	Variable	Coeff.
C	126 (4.8)	T^3	0.00002 (0.0)	ρ_1	-635. (27)
M_1	20.3 (0.77)	T^4	0.0000 (0.0)	ρ_1^2	5492 (285)
M_1^2	-2.28 (0.09)	N	1.51 (0.43)	ρ_1^3	-12829 (892)
M_1^3	0.092 (0.01)	N^2	-0.172 (0.03)	ρ_2	1592 (51)
M_1^4	-0.001 (0.00)	N^3	0.004 (0.00)	ρ_2^2	-16334 (981)
M_2	-1.85 (0.10)	N^4	-0.0000 (0.00)	ρ_2^2	73334 (7146)
M_2^2	0.094 (0.003)	$M_1 \cdot T$	0.010 (0.00)	ρ_2^4	-119671 (17403)
$M1 \cdot M2$	-0.085 (0.003)	$T \cdot N$	-0.003 (0.001)	$\rho_1 \cdot \rho_2$	-508 (44)
T	2.43 (0.06)	$M_1 \cdot N$	0.125 (0.004)		
T^2	-0.011 (0.00)	$M_2 \cdot N$	-0.013 (0.00)		
$R^2 = 0.943$ Number obs.: 10,000					

Table 1: Regression results for firm one. Dependent variable is 10,000 · Π_1 . Standard errors in parentheses.

The results for profit are reported in table 1, from which we can draw a number of conclusions.

- The regression confirms the finding that we discussed in other papers

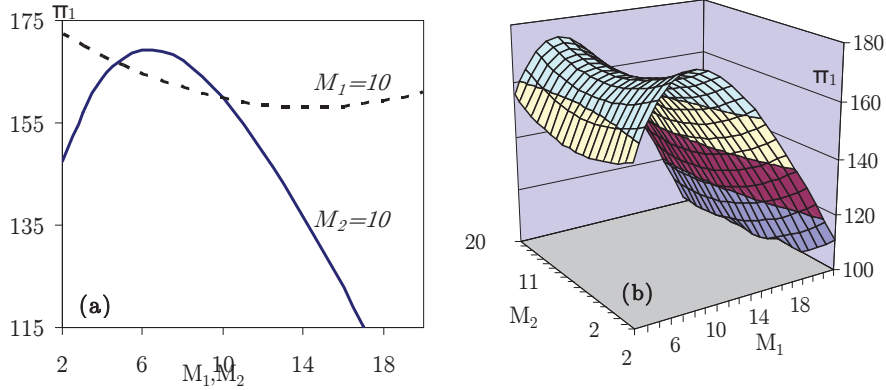


Figure 5: $\Pi_1(M_1, M_2)$. Bilateral relationship (quadrant *a*, where the continuous line is $\Pi_1 = f(M_1, 10)$ and the dotted line is $\Pi_1 = f(10, M_2)$) and joint effect (quadrant *b*).

(Barr and Saraceno, 2002, 2005) on the humped-shaped relationship between firm dimension and profits. The optimal (i.e., profit maximizing) size is determined by balancing speed and accuracy in learning. Figure 5 plots the coefficients of size (M_1 and M_2) from table 1. The continuous line of quadrant *a* shows that, at first, profit increases with own firm size, and then drops, showing that the optimal balance of speed and accuracy lies away from the extremes. The dotted line shows that increasing the rival's size has a modest u-shaped effect on a firm's profits. A rival that is very small or very large confers a competitive advantage on the firm because it is either relatively inaccurate (when small) or too slow (when large). However, the overwhelming importance of own size in determining profits also emerges when we look at the joint effect (quadrant *b*), where the effect of M_2 and of the interaction term is relatively small.

- Environmental complexity N has a negative and nonlinear effect on profits, while increasing the time to learn (T) has a positive effect. Both findings conform to intuition.
- The relationship between profit and niceness is shown in figure 6, where we plot the regression coefficients of ρ_1 and ρ_2 . Quadrant *a* shows that extreme values of own willingness to cooperate, are more profitable:




Figure 6: Profit of firm one against the niceness parameters ρ_1 and ρ_2 , bilateral relationships (quadrants *a* and *b*, where $\rho_2 = 0$ and $\rho_1 = 0$, respectively), and joint effect.

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, higher willingness to cooperate from the opponent (larger ρ_2 , quadrant *b*) yields larger profits for firm 1. When we allow ρ_1 and ρ_2 to change contemporaneously, we can confirm this finding.

We now turn to the other dependent variable, the one specific to this paper—the degree of cooperation; from the regression coefficients (table 2) we can conclude that:

- As expected, larger values of ρ_1 and ρ_2 increase the cooperation rate (fig. 7, right quadrant). This is especially true for own willingness ρ_1 .
- 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 (fig. 7 left quadrant). Concerning a firm’s own size (the continuous line), the relationship is negative, albeit slightly so; larger firms

Dependent Variable: $10,000 \cdot C_1$					
Variable	Coeff.	Variable	Coeff.	Variable	Coeff.
C	3613 (55)	T^2	-0.003 (0.00)	ρ_2	26970 (1237)
M_1	-9.4 (2.5)	N	-15.5 (0.82)	ρ_2^2	-262625 (25083)
M_1^2	0.22 (0.08)	$M_2 \cdot N$	0.844 (0.04)	ρ_2^3	1069786 (188273)
M_2	223 (17)	$T \cdot N$	-0.0271 (0.00)	ρ_2^4	-1537810 (467083)
M_2^2	-22.6 (2.3)	ρ_1	84001 (1235)	$M_1 \cdot \rho_1$	29.7 (8.9)
M_2^3	0.934 (0.13)	ρ_1^2	-700449 (25094)		
M_2^4	-0.015 (0.00)	ρ_1^3	2996843 (188970)		
T	3.75 (0.14)	ρ_1^4	-5197243 (469096)		
$R^2 = 0.925$ Number obs.: 10,000					

Table 2: Regression results. Dependent variable is 10,000*cooperation rate of firm one. Standard errors in parentheses.

tend to cooperate less often. Interestingly, a rival's size has a relatively strong influence on a firm's cooperation rate, which furthermore is hump shaped. Notice the similarity with the profit function, which suggests a correlation with learning performance: the better a rival is able to learn the environment, the more likely a firm is to cooperate, since its rivals actions are less noisy.

To summarize, the regressions confirm the robustness of the results of section 6. This model conforms to previous findings (Barr and Saraceno, 2002; 2005) on the relationship between firm profitability, size and the environment it faces. Environmental complexity negatively affects profits, and the tradeoff between speed and accuracy emerges in this setting as well, giving a hump shape relationship between firm size and profit. We also showed that firms take advantage of their opponent's willingness to cooperate, and that intermediate levels of niceness yield worse results than extreme (harsh or soft) attitudes. The regressions also show that cooperation is hampered by more complex environments, and, of course, by lower willingness to cooperate. In addition, our regression analysis sheds light on the relationship




Figure 7: Left: C_1 as a function of M_1 (continuous) and M_2 (dashed). Right: C_1 as a function of ρ_1 (continuous) and ρ_2 (dashed).

between firms size and cooperation.

8 Network Equilibria

In this section we explore the concept of *network equilibrium* with regard to the two important factors that affect performance: network size and the level of niceness. We define a network equilibrium (NE) as a Nash equilibrium in M_j and ρ_j , i.e., a quadruple $\{M_1^*, \rho_1^*, M_2^*, \rho_2^*\}$ such that neither firm finds it profitable to change its network size or niceness given its rival's choice.

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

Notice that we do not endogenize the choices of M_j and ρ_j . Rather, we conduct a kind of comparative statics exercise, whereby we study the NEs for given environmental conditions. To run this experiment, we fix the number of iterations at $T = 250$. We have each firm play against a rival of all possible size and niceness values. For example, we begin by setting $M_1 = 2$, and $\rho_1 = 0$. Then firm one, with its size and niceness fixed, plays against a rival of all different sizes $M_2 \in \{2, 3, \dots, 20\}$, and niceness values $\rho_2 \in \{0, 0.125, \dots, 0.2\}$. Then we repeat the exercise for $M_1 = 3$, and $\rho_1 = 0$, and so on. For each set of M 's and ρ 's we generate 50 runs and take averages to smooth out fluctuations from each run. We then determine, for each complexity value and for each firm, the network size and niceness value that is the most profitable against




Figure 8: Optimal profit and firm dimension, industry averages.

its rival. That is, we search for each firm’s “best response” network size and niceness. The equilibrium is then given by the sizes and niceness values which are a “best response” for both firms.

We repeated these runs for environmental complexity values of $N = 5, 10, \dots, 30$, and calculated the NEs that emerged (for each number of inputs there were always more than one equilibrium). Finally, we averaged the total number of managers of the two firms for each equilibrium to obtain an equilibrium “industry” size and averaged the ρ 's for an average “industry” niceness. This gave us a data set of 96 NEs (an average of 16 per complexity level).

Profits and Network Size Figure 8 shows that higher complexity requires larger firm size in equilibrium, which is necessary to better learn its rival’s actions and the environment. Furthermore, the increased difficulty of learning is associated with lower profits because the additional benefit of larger firm size does not fully offset the marginal increase in errors that firms make due to increased environmental complexity.

Figure 9 shows, however, that a clear pattern between complexity and cooperation does not appear. The equilibrium average ρ shows a slight in-

crease in complexity, going from 0.078 to 0.086. Interestingly, the degree of cooperation does not seem to be related to the number of inputs. In fact, even if as shown in table 2, a firm’s own niceness has a large impact on its cooperation rate, the effect of a rival’s niceness is more modest. Thus, the effect of the interaction of the two parameters on total industry behavior will not be strong. In addition, on one hand complexity is likely to lower cooperation, but on the other hand, it will also generate larger firm sizes in order for firms to better learn the environment, which is likely to increase cooperation. As a consequence larger firm sizes and more complexity will cancel each other out.




Figure 9: Equilibrium cooperation rates and niceness values. Industry averages.

9 Conclusion

This paper has presented a model of firm learning and cooperation. We investigate the prospects for cooperation when firms must learn to map environmental signals to changing demand and on their rival’s output decision. We demonstrate that increased environmental complexity is associated with

lower rates of cooperation and lower profitability. We show that 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 own size increases profits up to a point and then profits decrease. This tradeoff is due to the fact that adding agents improves accuracy but slows learning. We also see a non-linear relationship between the firm size and its cooperation rate; in particular a humped-shaped relationship between a rival’s size and a firm’s cooperation rates. Finally, we have shown that in a network equilibrium, increasing complexity is associated with larger average firm size but, given the interplay of different and conflicting factors, has little effect on cooperation.

Clearly our model of the firm is highly stylized, but the goal here is to show how the firm-as-learning-network can affect and be affected by the competitive environment. The aim of our work is to illustrate how firm organization can have a strategic impact on both firm’s profits, and on its ability to send and receive competitive information. Large organizations emerge precisely because no single agent can do even a small fraction of the work needed to supply goods to the marketplace and to maximize profits. A firm’s environment is often quite complex and thus some agents must be devoted to learning the nature of this environment. Our model illustrates that multi-agent learning can have an impact on not only the firm, but also on industry wide outcomes. In other words organizational structure can affect tacit communication between organizations, when they are unable to communicate directly.

We believe that economics has only begun to think about the firm as a collection of several agents; yet how a firm organizes itself can have important consequences for competition. The large body of literature within management studies on firm reorganization and restructuring is evidence of the causes and effects of firm organization on performance and hence competition. (See, for example, *Strategic Management Journal*, special issue on corporate restructuring, 1993.)

The resource allocation and organizational problems of real firms are quite complex. And perhaps because of this complexity, economists have shied away from trying to present models of the firm as an organization of agents. Our work is but a first attempt to model this complexity. The next steps in this process would be the attempt introduce some more realistic agents and firm behaviors, such as including heterogeneity in strategies and endogenizing

the hierarchical structure of the network.

References

- Axelrod, R., 1984. *The Evolution of Cooperation*. New York: Basic Books.
- Axelrod, R., 1997. *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*. Princeton: Princeton University Press.
- Axtell, R. L., Epstein, J. M., 1996. *Growing Artificial Societies: Social Science from the Bottom Up*. Cambridge: The MIT Press.
- Barr, J., Saraceno, F., 2005. Cournot competition, organization and learning. *Journal of Economic Dynamics and Control*, 29, 277-295.
- Barr, J., Saraceno, F. 2002. A computational theory of the firm. *Journal of Economic Behavior and Organization*, 49, 345-361.
- Camerer, C. and Vepsäläinen, A., 1988. The economic efficiency of corporate culture. *Strategic Management Journal*, 9, 115-126.
- Carley, K. M., 1996. A Comparison of artificial and human organizations. *Journal of Economic Behavior and Organization*, 31, 175-191.
- Casson, M., 1991. *The Economics of Business Culture: Game Theory, Transaction Costs and Economic Performance*. Oxford: Calredon Press.
- Chang, M.-H., Harrington, J. E., 2006. Agent-based models of organizations. In: Leigh Tesfatsion, L., Judd, K. L. (Eds.). *Handbook of Computational Economics*, Vol. II. Amsterdam: North-Holland, 1272-1337.
- Cho, I.-K., 1994. Bounded rationality, neural networks and folk theorem in repeated games with discounting. *Economic Theory*, 4, 935-957.
- Croall, I. F., Mason, J. P., 1992. *Industrial Applications of Neural Networks: Project ANNIE Handbook*. New York: Springer-Verlag.
- Cyert, R.M., DeGroot, M.H., 1973. An analysis of cooperation and learning in a duopoly context. *American Economic Review*, 63, 24-37.
- DeCanio, S. J., Watkins, W. E., 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, E., Porter R., 1984. Noncooperative Collusion under Imperfect Price Competition. *Econometrica*, 52, 87-100.
- Hannon, M. T., Freeman, J. 1989. *Organizational Ecology*. Cambridge: Harvard University Press.
- Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., Gintis, H., McElreath, R., 2001. In search of homo economicus: behavioral experiments in 15 small-scale societies. *American Economic Review*, 91(2), 73-78.
- 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*. Ann Arbor: University of Michigan Press.
- Kuan C.-M., White, H., 1994. Artificial neural networks: an econometric perspective. *Econometric Reviews*, 13, 1-91.
- Levenstein, M. C., Suslow, V.Y., 2006. What determines cartel success? *Journal of Economic Literature*, 44(1), 42-95.
- Miller, J. H., 1996. The coevolution of automata in the repeated prisoner's dilemma. *Journal of Economic Behavior and Organization*, 29, 87-112.
- Nelson, R. R. , Winter, S. G., 1982. *An Evolutionary Theory of Economic Change*. Cambridge: Belknap Press of Harvard University Press.
- Niyogi, P., 1998. *The Informational Complexity of Learning: Perspectives on Neural Networks and Generative Grammar*. Boston: Kluwer Academic Publishers.
- Radner, R., 1993. The organization of decentralized information processing. *Econometrica*, 61, 1109-1146.
- Rubinstein, A. R., 1986. Finite automata play the repeated prisoners dilemma. *Journal of Economic Theory*, 39, 83-96.
- Skapura, D. M., 1996. *Building Neural Networks*, New York: Addison-Wesley.

- Strategic Management Journal. 1993. 14, summer issue.
- Tesfatsion, L. 2002. Agent-based computational economics: growing economies from the bottom up. *Artificial Life*, 8(1), 55-82.
- Tirole, J., 1988. *The Theory of Industrial Organization*, Cambridge: The MIT Press.
- Verboven, F. 1997. Collusive behavior with heterogeneous firms. *Journal of Economic Behavior and Organization*, 33, 122-139.
- Viscusi, W. K., Vernon, J. M., Harrington, J. E., 2001. *Economics of Regulation and Antitrust*, 3rd ed. Cambridge: The MIT Press.
- 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.
- White, H., 1992. *Artificial Neural Networks: Approximation and Learning Theory*. Cambridge: Blackwell.
- Wooldridge, J. M., 2003. *Introductory Econometrics: A Modern Approach*, 2nd ed. Mason: Thompson South-Western .