Simulating the Ecology of Oligopoly Games with Genetic Algorithms*

Shu-Heng Chen
AI-ECON Research Group
Department of Economics
National Chengchi University
Taipei, Taiwan 11623
TEL: 886-2-29387308
FAX: 886-2-27386874
E-mail: chchen@cc.nccu.edu.tw

Chih-Chi Ni
AI-ECON Research Group
Department of Economics
National Chengchi University
Taipei, Taiwan 11623
TEL: 886-2-29387308
FAX: 886-2-29390344
E-mail: g2258503@grad.cc.nccu.edu.tw

Abstract

In this paper, the relation between the N-person IPD game and the N-person oligopoly game is rigorously addressed. Our analytical framework shows that due to the path-dependence of the payoff matrix of the oligopoly game, the two games in general are not close in spirit. We then further explore the significance of the path-dependence property to the rich ecology of oligopoly from an evolutionary perspective. More precisely, we simulated the evolution of a 3-person oligopoly game, and showed that the rich ecology of oligopoly can be exhibited by modelling the adaptive behaviour of oligopolists with genetic algorithms. The emergent behaviour of oligopolists are presented and analyzed. We indicate how the path-dependence nature may shed light on the phenotypes and genotypes coming into existence.

Keywords: Oligopoly, Market-Share First, Predatory Pricing, Genetic Algorithms, State-Dependent Markov Chain, Coevolution.

---

*This is a revised and extended version of the paper "Using Genetic Algorithms to Simulate the Evolution of an Oligopoly Game" presented at *The Second Asia-Pacific Conference on Simulated Evolution and Learning* in Canberra, Australia, 24-27 November, 1998. The authors wish to thank the two anonymous referees for their helpful comments. Research support from NSC grant NSC. 86-2415-H-004-022 is also gratefully acknowledged.

1 Motivation and Introduction

In industrial economics, modelling a market consisting of only a few firms, i.e., oligopolistic industry, is a highly indeterminate subject. In this area, economists are at variance with each other even on the most basic issue, i.e., *how price is determined?* [1] Since the indeterminacy of this subject may arise from the *perplexing interdependent relations and interactions among firms*, the relevance of *game theory* to oligopoly theory seems to be quite obvious. (Fudenberg and Tirole, 1989) In fact, a game known as the *iterated prisoner's dilemma* (**IPD**) *game* is frequently cited as an effective abstraction of the oligopoly pricing problem in many textbooks. Rephrasing the oligopoly pricing problem in the game context, many economists consider the oligopoly game and the IPD game close in spirit. Implicitly, it is assumed that players in the N-person oligopoly game are facing a pretty similar situation as players in the N-person oligopoly game.

Recently, the n-person IPD game was studied in Yao and Darwen (1994). Using *genetic algorithms* (GAs), they showed that *cooperation can still be evolved in a large group, but that it is more difficult to evolve cooperation as the group size increases.* Considering this result as a guideline for the oligopoly pricing problem, then what the n-person IPD game tells us is that *when the number of oligopolists is small, say 3, it is very likely to see the emergence of collusive pricing (cooperation).* Real data, however, usually shows that even in a three-oligopolist industry the observed pricing pattern is not that simple. (Midgely, Marks and Cooper, 1997)[2]

- First, while collusive pricing is frequently observed, it is continually interrupted by the occurrence of predatory pricing (price wars).
- Second, it is not always true that oligopolists are *collectively* charging high prices (collusive pricing) or low prices (price wars). In fact, a dispersion of prices can persistently exist, i.e., some firms are charging high prices, whilst others are charging low prices.
- Third, the firms who charge low prices may switch to high prices in a later stage, and vice versa.

These features may be best summarized by a quotation from Scherer and Ross (1990), a leading textbook in industrial economics.

> Casual observation suggests that in oligopoly *virtually anything can happen*. Some industries–cigarettes and breakfast cereals come readily to mind–succeed in maintaining prices well above production costs for years. Others, despite conditions that would appear at first glance to encourage cooperative behaviour, gravitate toward price warfare." (ibid, p.199; Italics added)

Nonetheless, "*virtually anything can happen*" is not the property which one may experience from a 3-person IPD game (See Yao and Darwen, ibid, Figure 5), and this gives us two questions.

[1] For a survey of the oligopoly literature, see Shapiro (1989).

[2] The overall patterns of prices and sales for the three major brands of coffee, Maxwell House Regular, Folgers, and Chock Full O'Nuts, can be found in Midgely, Marks and Cooper (1997).

- Is the N-person oligopoly game *close in spirit* to the N-person IPD game?
- If not, can we replicate the rich ecology of the N-person oligopoly game by *just* simulating the evolution of the oligopoly game?

The contribution of this paper is *two-fold*. First, contrary to some people's presumption, we show that the N-person oligopoly game is in general *not* close in spirit to the N-person IPD game. This may not be a revelation. *However, what was not seen in the past is a rigorous analysis of the argument, and this paper filled the gap.* In Section 2, we propose a very simple oligopoly game with 3 oligopolists. In this game, the payoff matrix is determined by *the market share dynamics*, which is characterized by a *time-variant state-dependent Markov transition matrix*. This framework enables us to see an important property of the game, i.e., *the path-dependence of the game*. Through this property, we can see that while these two games in general are not close in spirit, *there does exist a trivial path* on which these two games are *effectively* the same. Therefore, we believe that our analysis provides a general picture of the relation between the two games.

Given the *path-dependence* property of the oligopoly game, we further explore its significance to the ecology of the oligopoly game from an *evolutionary* perspective. By that we mean to account for the rich ecology of oligopoly *solely* from an evolutionary standpoint. In other words, unlike many conventional studies in this area, we do not attempt to build up an explanation by introducing *outside factors*, such as economic fluctuations, structural changes and institutional arrangements. (Green and Porter,1984; Abreu, Pearce and Stacchetti, 1986) Instead, we are asking: *other things being equal, can we still have a rich ecology of oligopoly*? We consider this effort a search for a more *fundamental cause*. This comes to the second contribution of this paper.

In section 3, we illustrate the use of the genetic algorithm (GAs) to model the adaptive behaviour of oligopolists. The application of GAs to the oligopoly game is nothing new. Midgely, Marks and Cooper (1997) pioneered this line of this research. While, we follows the ideas employed in their paper in many aspects, there is an important distinction: what Midgely et al. did was to use *historical market data* to breed GA-based oligopolists for the purpose of developing competitive marketing strategies; our paper has a different focus. We are not studying how GA-based oligopolists can compete with *real managers*, which is more like an application in the machine-learning style and Midgely et al. already did an excellent job on it. What concerns us instead is to use the GA to make histories of the oligopoly game on its own and see how rich these ecologies can be. Therefore, it is enough to just let our GA-based oligopolists learn from their own experiences and acquire expertise without being exposed to real data.

As what we shall see in Section 4, what were demonstrated from our simulations is a very rich ecology of oligopoly. Our analysis of the simulation results shows how this rich ecology can be related to the path-dependent property of the oloigopoly game. Some interesting patterns such as the *market-share first* strtegy, *nonagreesion argeement*, *unbalanced market power* and *death of firms* are also discussed. Concluding remarks are given in Section 5.

2 The Analytical Model

For simplicity, an oligopoly industry is assumed to consist of three firms, say $i = 1, 2, 3$. At each period, a firm can either charge a high price P_h or a low price P_l. Let a_i^t be the action taken by firm i at time t. $a_i^t = 1$ if the firm i charges P_h and $a_i^t = 0$ if it charges P_l. To simplify notations, let S_t denote the row vector (a_1^t, a_2^t, a_3^t). To characterize the price competition among firms, the *market share dynamics* of these three firms are summarized by the following *time-variant state-dependent Markov transition matrix*,

$$M_t = \begin{bmatrix} m_{11}^t & m_{12}^t & m_{13}^t \\ m_{21}^t & m_{22}^t & m_{23}^t \\ m_{31}^t & m_{32}^t & m_{33}^t \end{bmatrix} \tag{1}$$

where m_{ij}^t, the transition probability from state i to state j, denotes the proportion of the customers of firm i switching to firm j at time period t. Let n_i^t (i=1,2,3) be the number of customers of firm i at time period t, and N_t the row vector $[n_1^t, n_2^t, n_3^t]$. Without loss of generality, we assume that each consumer will purchase only one unit of the commodity. In this case, N_t is also the vector of quantities consumed. With N_t and M_t, the customers of each firm at period $t+1$ can be updated by:

$$N_{t+1} = N_t M_t \tag{2}$$

To see the effect of price competition on the market share dynamics, the transition probabilities m_{ij}^t are assumed to be dependent on the pricing strategy vector S_t. If three firms charge the same price, then M_t is an *identity matrix*. Furthermore, if firm i charges P_h, then it will lose $\frac{\alpha}{2} \times 100$ percent of its consumers each to firms j and k, who charge P_l. Furthermore, if firms i and j charge P_h, then they each will lose $\alpha \times 100$ percent of their consumers to firm k, who charges P_l. These assumptions can be summarized by the following transition matrices.

$$M_t(1,1,1) = M_t(0,0,0) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

$$M_t(1,0,0) = \begin{bmatrix} 1-\alpha & \frac{\alpha}{2} & \frac{\alpha}{2} \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad M_t(0,1,0) = \begin{bmatrix} 1 & 0 & 0 \\ \frac{\alpha}{2} & 1-\alpha & \frac{\alpha}{2} \\ 0 & 0 & 1 \end{bmatrix},$$

$$M_t(0,0,1) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ \frac{\alpha}{2} & \frac{\alpha}{2} & 1-\alpha \end{bmatrix}, \quad M_t(1,1,0) = \begin{bmatrix} 1-\alpha & 0 & \alpha \\ 0 & 1-\alpha & \alpha \\ 0 & 0 & 1 \end{bmatrix},$$

$$M_t(1,0,1) = \begin{bmatrix} 1-\alpha & \alpha & 0 \\ 0 & 1 & 0 \\ 0 & \alpha & 1-\alpha \end{bmatrix}, \quad M_t(0,1,1) = \begin{bmatrix} 1 & 0 & 0 \\ \alpha & 1-\alpha & 0 \\ \alpha & 0 & 1-\alpha \end{bmatrix}$$

Given these state-dependent transition matrices, Equation (2) can be rewritten as:

$$N_{t+1} = N_t M_t(S_t), \tag{3}$$

where $S_t = (a_1^t, a_2^t, a_3^t)$ and $a_i^t \in \{0, 1\}$. Equation (3) summarizes the *intra*-industry competition given a number of customers $n_t = \sum_{i=1}^{3} n_i^t$.

Given Equation (3), the objective of oligopolists is to maximize their profits or the present value of the firm, and the profits for a single period is given by Equation (4).

$$\pi_i^s = (P_i^s - C)n_i^s \tag{4}$$

where P_i^s is the price charged by firm i at period s, n_i^s number of customers, and C fixed unit cost. n_i^s can be obtained from Equation (3).

Given this simple framework of the oligopoly game, we would like to know, to what extent, this simple oligopoly game can be related to the n-person IPD game. More precisely, *is the oligopoly game defined above necessarily an n-person IPD game?* To answer this question, we have to work out the payoff matrix used to define an N-person IPD game. (Yao and Darwen, 1994) However, due to the dynamics of market shares, the payoff matrix is in general *not static*. We, therefore, start our analysis from the first-round of the oligopoly game. Suppose that each round of the oligopoly game consists of r iterations of the game, and that "cooperate" (C) means "charging high prices for all r periods" and "defect" (D) means "charging low prices for all r periods".

We can now work out the first-round *payoff matrix* employed by Yao and Darwen (1994). In our case (3 oligopolists), there are six elements in the payoff matrix, namely C_i and D_i ($i = 0, 1, 2$). Here, C_i (D_i) denotes the payoff for a specific player who plays C (D) when there are i players acting cooperatively. Without losing generality, let us assume that $n_1^1 = n_2^1 = n_3^1 = 1$; then C_i and D_i can be computed from the following four equations:

$$\begin{bmatrix} D_2 & C_1 & C_1 \end{bmatrix} = \sum_{t=1}^{r} \begin{bmatrix} 1 & 1 & 1 \end{bmatrix} [M(0,1,1)]^t \begin{bmatrix} P_L - C & 0 & 0 \\ 0 & P_H - C & 0 \\ & 0 & P_H - C \end{bmatrix} \tag{5}$$

$$\begin{bmatrix} C_2 & C_2 & C_2 \end{bmatrix} = \sum_{t=1}^{r} \begin{bmatrix} 1 & 1 & 1 \end{bmatrix} [M(1,1,1)]^t \begin{bmatrix} P_H - C & 0 & 0 \\ 0 & P_H - C & 0 \\ & 0 & P_H - C \end{bmatrix} \tag{6}$$

$$\begin{bmatrix} D_1 & D_1 & C_0 \end{bmatrix} = \sum_{t=1}^{r} \begin{bmatrix} 1 & 1 & 1 \end{bmatrix} [M(0,0,1)]^t \begin{bmatrix} P_L - C & 0 & 0 \\ 0 & P_L - C & 0 \\ & 0 & P_H - C \end{bmatrix} \tag{7}$$

and

$$\begin{bmatrix} D_0 & D_0 & D_0 \end{bmatrix} = \sum_{t=1}^{r} \begin{bmatrix} 1 & 1 & 1 \end{bmatrix} [M(0,0,0)]^t, \begin{bmatrix} P_L - C & 0 & 0 \\ 0 & P_L - C & 0 \\ & 0 & P_L - C \end{bmatrix} \tag{8}$$

where

$$
\begin{aligned}
[M(0,1,1)]^r &= \begin{bmatrix} 1 & 0 & 0 \\ \alpha & 1-\alpha & 0 \\ \alpha & 0 & 1-\alpha \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ \alpha & 1-\alpha & 0 \\ \alpha & 0 & 1-\alpha \end{bmatrix}^{r-1} \\
&= \begin{bmatrix} 1 & 0 & 0 \\ \alpha+\alpha(1-\alpha) & (1-\alpha)^2 & 0 \\ \alpha+\alpha(1-\alpha) & 0 & (1-\alpha)^2 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ \alpha & 1-\alpha & 0 \\ \alpha & 0 & 1-\alpha \end{bmatrix}^{r-2} \\
&= \begin{bmatrix} 1 & 0 & 0 \\ \sum_{j=0}^{r-1} \alpha(1-\alpha)^j & (1-\alpha)^r & 0 \\ \sum_{j=0}^{r-1} \alpha(1-\alpha)^j & 0 & (1-\alpha)^r \end{bmatrix},
\end{aligned}
$$

$$
\begin{aligned}
[M(0,0,1)]^r &= \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ \frac{\alpha}{2} & \frac{\alpha}{2} & 1-\alpha \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ \frac{\alpha}{2} & \frac{\alpha}{2} & 1-\alpha \end{bmatrix}^{r-1} \\
&= \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ \frac{\alpha+\alpha(1-\alpha)}{2} & \frac{\alpha+\alpha(1-\alpha)}{2} & (1-\alpha)^2 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ \frac{\alpha}{2} & \frac{\alpha}{2} & 1-\alpha \end{bmatrix}^{r-2} \\
&= \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ \frac{\sum_{j=0}^{r-1} \alpha(1-\alpha)^j}{2} & \frac{\sum_{j=0}^{r-1} \alpha(1-\alpha)^j}{2} & (1-\alpha)^r \end{bmatrix},
\end{aligned}
$$

and

$$
[M(1,1,1)]^r = [M(0,0,0)]^r = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}^r .
$$

A few steps of computation will show:

$$
\begin{aligned}
\begin{bmatrix} D_2 & C_1 & C_1 \end{bmatrix}' &= \begin{bmatrix} (P_L - C)(r + 2\sum_{t=1}^{r}\sum_{s=0}^{t-1} \alpha(1-\alpha)^s) \\ (P_H - C)(\sum_{t=1}^{r}(1-\alpha)^t) \\ (P_H - C)(\sum_{t=1}^{r}(1-\alpha)^t) \end{bmatrix} \\
&= \begin{bmatrix} (P_L - C)[3r - 2\frac{(1-\alpha)-(1-\alpha)^{r+1}}{\alpha}] \\ (P_H - C)[\frac{(1-\alpha)-(1-\alpha)^{r+1}}{\alpha}] \\ (P_H - C)[\frac{(1-\alpha)-(1-\alpha)^{r+1}}{\alpha}] \end{bmatrix},
\end{aligned} \tag{9}
$$

$$
\begin{aligned}
\begin{bmatrix} D_1 & D_1 & C_0 \end{bmatrix}' &= \begin{bmatrix} (P_L - C)[r + \frac{1}{2}\sum_{t=1}^{r}\sum_{s=0}^{t-1} \alpha(1-\alpha)^s] \\ (P_L - C)[r + \frac{1}{2}\sum_{t=1}^{r}\sum_{s=0}^{t-1} \alpha(1-\alpha)^s] \\ (P_H - C)(\sum_{t=1}^{r}(1-\alpha)^t) \end{bmatrix} \\
&= \begin{bmatrix} (P_L - C)[r + \frac{1}{2}r - \frac{1}{2}\sum_{j=1}^{r}(1-\alpha)^j] \\ (P_L - C)[r + \frac{1}{2}r - \frac{1}{2}\sum_{j=1}^{r}(1-\alpha)^j] \\ (P_H - C)(\sum_{t=1}^{r}(1-\alpha)^t) \end{bmatrix},
\end{aligned} \tag{10}
$$

Table 1: Parameters and Payoffs: The First Round of the Game

Set	P_H	P_L	C	α	r	D_2	D_1	D_0	C_2	C_1	C_0	$(D_2+C_1)/2$	$(D_1+C_0)/2$
1	1.4	1.2	1	0.2	8	3.47	2.07	1.6	3.2	1.33	1.33	2.4	1.70
2	1.4	1.2	1	0.2	25	13.40	7.10	5	10	1.60	1.60	7.50	4.35

$$\begin{bmatrix} C_2 & C_2 & C_2 \end{bmatrix}' = \begin{bmatrix} (P_H-C)r \\ (P_H-C)r \\ (P_H-C)r \end{bmatrix}, \begin{bmatrix} D_0 & D_0 & D_0 \end{bmatrix}' = \begin{bmatrix} (P_L-C)r \\ (P_L-C)r \\ (P_L-C)r \end{bmatrix}. \tag{11}$$

Based on the derived payoff vector $(D_2, D_1, D_0, C_2, C_1, C_0)$, we can decide whether the oligopoly game is an n-person IPD game by checking the following criteria (Yao and Darwen, 1994: Figure 2):

- (1) $D_2 > C_2$, (2) $D_1 > C_1$, and (3) $D_0 > C_0$.
- (4) $D_2 > D_1 > D_0$, and (5) $C_2 > C_1 > C_0$.
- (6) $C_2 > \frac{D_2+C_1}{2}$, and (7) $C_1 > \frac{D_1+C_0}{2}$.

The first five conditions feature the conflict between two forces, namely, the *temptation to defect* and the *fear of retaliation*. The last two conditions are somewhat *tricky*. They exclude the possibility of the other type of *cooperation*, i.e., *false defection*. In the prisoner's dilemma game, one prisoner can be willingly betrayed, allowing them to reap the reward. He will then have a share of the reward as a compensation for his sacrifice. Like cooperation, false defection requires a delicate design and is intelligent behaviour. It is interesting to note that, in reality oiligopolists cut their prices *in turn*. Superficially, these actions can be interpreted as a result of competition, but, in effect, they are another type of collusive pricing when the game is not bounded by the last two conditions. Since the failure to meet the last two conditions implies another type of intelligent behaviour, it is useful to take a notice of this emergent intelligence in our simulation to be discussed later.

By Equations (9)-(11), the payoff vector is a function of P_H, P_L, C, r and α. It is not difficult to see that, in general, not all of these conditions can be satisfied. For example, in Table 1, two sets of parameters and their associated payoffs are given. The conditions which can be satisfied by these two sets of parameters are summarized in Table 2. Among them, Condition 7 is strictly violated in both cases. Nevertheless, since the first five conditions are satisfied, the oligopoly game shares the essential ingredients of the N-person IPD game, namely, the temptation to defect and the fear of retaliation.

So far, we have only worked out the payoff vector of the first round of the game. The payoff vector of the second round, and the rounds after, is a little intriguing. Due to the dynamics of market shares, the payoff vector is not independent of what happened in the first round. In other words, the payoff vector, like the market-share dynamics, is also *time-variant* and *path-dependent*. To see this, it is helpful to work out the second-round payoff vector too. Since there are six nonredundant histories (paths) in the first round, and with each there is a follow-up

Table 2: Parameter Sets and Testing Results

Inequality	Set 1	Set 2
1. $D_2 > C_2$	$>$	$>$
2. $D_1 > C_1$	$>$	$>$
3. $D_0 > C_0$	$>$	$>$
4. $D_2 > D_1 > D_0$	$>,>$	$>,>$
5. $C_2 > C_1 > C_0$	$>,=$	$>,=$
6. $C_2 > 0.5(D_2 + C_1)$	$>$	$>$
7. $C_1 > 0.5(D_1 + C_0)$	$<$	$<$

The sign $>$ in columns 2 and 3 means the condition is satisfied. Other signs mean the condition is weakly violated ($=$) or strongly violated ($<$).

payoff vector, the second-round payoffs of the game can be represented by a 6 by 6 payoff table as shown in Table 3.

Table 3 exhibits the payoff matrix of Parameter Sets 1 and 2. Clearly, this matrix is much more complicated than the one in the first round of the game. It differs from the first-round payoff matrix in three ways. First of all, in the first round of the game the payoff is symmetric, while in the second, it depends. For example, if in the first round, three firms all charge the high price or low price, for that matters, then in the second round, the initial condition for them is the same, and hence their payoff vectors refer to the same row led by C2 or D0 in Table 3. In this case, the symmetry of payoffs remains unchanged. However, suppose in the first round, two firms charge the high price, and one firm charges the low price; then, in the second round, the payoff vector to the firm who charges the high price is the one led by C1, while the one to the firm who charges the low price is the one led by D2. According to Table 3, there two rows are not identical; consequently, the symmetric property does not hold. Therefore, *in the oligopoly game, whether the symmetry property will hold depends on the path of the market dynamics.*

Second, in addition to asymmetry of the payoff vectors, it is interesting to notice that in some cases, the payoff is *not unique*. For example, in Table 3, if the initial state for the player is D1, and the current state is also D1, then the payoff can be either 2.34 or 2.93. These non-unique outcomes result from the fact that after the first round of the game, each player may have different market shares. Therefore, it is not just *the number of cooperators* (among the remaining n-1 players) that matters but, most of all, *who* cooperates. If a firm with a larger market share cooperates, then "defection" can be more profitable because of the proportion of the market one could seize.

Last, the reversal of the payoff inequality. Based on Table 3, we checked the payoff inequality under each initial condition, and the results are summarized in Table 4. From Table 4, we find that, depending on the initial conditions, the first and the sixth inequality can be reversed (see the sign in brackets). More precisely, when the initial condition is D2 or D1, instead of "$D2 > C2$", we have "$D2 < C2$", and when the initial condition is C1, "$C2 > (D2 + C1)/2$". The reversal of the first condition is *critical*, because if two players choose C, then the dominant

Table 3: Parameters and Payoffs: The Second Round of the Game

Parameter Set 1:								
Round 1/Round 2	D2	D1	D0	C2	C1	C0	$(D_2+C_1)/2$	$(D_1+C_0)/2$
D2	4.58	4.34	4.26	8.53	3.55	3.55	4.06	3.94
D1	3.75	2.34 2.93	2.27	4.53	1.89	1.89	2.82	2.11 2.41
D0	3.47	2.07	1.60	3.20	1.33	1.33	2.40	1.70
C2	3.47	2.07	1.60	3.20	1.33	1.33	2.40	1.70
C1	2.91	0.35 1.51	0.27	0.54	0.22	0.22	1.57	0.29 0.87
C0	2.91	0.93	0.27	0.54	0.22	0.22	1.57	0.58
Parameter Set 2:								
Round 1/Round 2	D2	D1	D0	C2	C1	C0	$(D_2+C_1)/2$	$(D_1+C_0)/2$
D2	14.99	14.97	14.96	29.92	4.77	4.77	9.88	9.87
D1	13.80	7.50 10.64	7.49	14.98	2.39	2.39	8.10	6.51 4.94
D0	13.41	7.10	5.00	10.00	1.59	1.59	7.50	4.35
C2	13.41	7.10	5.00	10.00	1.59	1.59	7.50	4.34
C1	12.61	0.03 6.31	0.02	0.04	0.01	0.01	6.31	0.02 3.16
C0	12.61	3.17	0.02	0.04	0.01	0.01	6.31	1.59

option is also C rather than D. If this can happen, then the oligopoly game is *essentially* not an IPD game. The reversal of the sixth condition coupled with the original violation of the seventh condition is far from minor because it defines another highly intelligent cooperative behaviour, i.e., "*false defection*" as discussed above.

In sum, the oligopoly game is not an N-person IPD game. Nevertheless, it is related to the N-person IPD game in a subtle way. In particular, among innumerable paths of the oligopoly game, there are many which are effectively equivalent to an N-person IPD game. In other words, the equivalence of the oligopoly game and the N-person IPD game is *path dependent*. However, considering learning a stochastic selection process, we cannot restrict our players' evolution only to those specific paths. Hence, the simulation results obtained from the N-person IPD game may not be applicable to the oligopoly game. For example, the *probability of the emergence of collusive pricing* of the 3-person oligopoly game may be quite different from that of a 3-person IPD game. It is, therefore, interesing to know whether the *path-dependence property* of the oligopoly game can generate more complex emergent behaviour than the N-person IPD game, and if so, *what they are* and *what they mean*.

Table 4: Parameter Sets and Testing Results

Inequality/Initial Condition	D2	D1	D0	C2	C1	C0
1. $D_2 > C_2$	[$<$]	[$<$]	$>$	$>$	$>$	$>$
2. $D_1 > C_1$	$>$	$>$	$>$	$>$	$>$	$>$
3. $D_0 > C_0$	$>$	$>$	$>$	$>$	$>$	$>$
4. $D_2 > D_1 > D_0$	$>,>$	$>,>$	$>,>$	$>,>$	$>,>$	$>,>$
5. $C_2 > C_1 > C_0$	$>,=$	$>,=$	$>,=$	$>,=$	$>,=$	$>,=$
6. $C_2 > 0.5(D_2 + C_1)$	$>$	$>$	$>$	$>$	[$<$]	$<$
7. $C_1 > 0.5(D_1 + C_0)$	$<$	$<$	$<$	$<$	$<$	$<$

The sign $>$ in columns 2-7 means the condition is satisfied. Other signs means the condition is weakly violated ($=$) or strongly violated ($<$). Signs in the brackets refer to the reversals of the payoff inequality in the second round of the game.

3 Modeling the Adaptive Behavior of Oligopolists with GAs

In this study, we shall simulate the oligopoly game by using *genetic algorithms*. The basic idea is to simulate the dynamics of the oligopoly game as a result of a sequence of *interactions* among the local shops owned by different oligopolists (chain stores). Based on her pricing strategy, each shop interacted locally with other shops' pricing. The price strategy of each shop was represented by a *binary string*, and hence the pricing strategies of all the shops owned by the same oligopolist were nothing but *a population of binary strings*. The genetic algorithm was then applied to mimicking the evolution of a collection of pricing strategies by evolving the population of binary strings.

Formally, the pricing strategy ϕ is a mapping:

$$\phi : \Omega \longrightarrow \{0, 1\}, \tag{12}$$

where Ω is the collection of all histories of $\{S_j\}_{j=1}^{t-1}$. However, this general version is difficult to be coded by genetic algorithms since the memory size required is infinite. Following Midgley et al (1997), we consider a special class of pricing strategy ψ,

$$\psi : \Omega_k \longrightarrow \{0, 1\}, \tag{13}$$

where Ω_k is the collection of all $\{S_{t-j}\}_{j=1}^{k}$. By this simplification, the oligopolist's memory is assumed to be *finite*.

Since each firm can only take two kinds of actions and there are three firms, we have 2^3 possible states in each period and 2^{3k} possible states in Ω_k. Therefore, to encode a pricing strategy ψ in Ω_k, we need a binary string with length 2^{3k}. Clearly, the length of the string increases exponentially with k. While, potentially, different choices of k may lead to quite different sets of strategies (Beaufils et al., 1998), the issue concerns us is the smallest value of k which can reasonably replicate the price dynamics of the oligopoly industry, and as we shall

Table 5: The Parameters of the GA-based Oligopoly Game

Memory size (k)	1
Number of oligopolists (chain stores)	3
Population size (l) (# of shops in each chain)	30
Number of periods in a single play (r)	8 (25)
Selection Scheme	Roulette-wheel selection
fitness function	Profits (π)
Number of generations evolved (Gen)	250 (80)
Number of periods (T)	2000
Crossover Style	One-Point Crossover
Crossover rate	0.8
Mutation rate	0.001
Immigration rate	0.001

see later, setting k to equal 1 is good enough to achieve this goal. [3] In Ω_1, ψ can be coded with an 8-bit string. For example, an 8-bit string $b_1 b_2 ... b_8$ means that if state j ($j = 1, 2, ..., 8$) occurs, the oligopolist will take action b_j ($b_j = 0, 1$).[4] The eight states are ordered as in the following sequence.

$$\underbrace{000}_{1}\underbrace{001}_{2}\underbrace{010}_{3}\underbrace{011}_{4}\underbrace{100}_{5}\underbrace{101}_{6}\underbrace{110}_{7}\underbrace{111}_{8}$$

Each state is represented by a 3-bit string. From the left to the right, the first bit refers to the action taken by firm 1 in the last period, the second bit refers to the action taken by firm 2, and so on. For instance, state "5" encoded as "100" means that the firm charged a high price (cooperated), but firms 2 and 3 charged a low price (defected) in the last period. If $b_5 = 0$ for firm 1, then firm 1 will take a revenge by charging a low price at this period. Given the encoding scheme described above, oligopolists' adaptive behavior is implemented by genetic algorithms as follows.

- **Step 1:**
  In the initial generation, a population of ψ is randomly generated for three oligopolists. Call it Gen_i^1 ($i = 1, 2, 3$).
  $$Gen_i^1 = \{\psi_i^1, ..., \psi_i^q, ..., \psi_i^l\} \tag{14}$$
  where l is the size of population (the number of chromosomes).

[3] It remains to be seen whether a high value of k can significantly change the result. If this is the case, then we should seriously consider an economic interpretation of the parameter k. We are currently conducting this line of research.

[4] As Yao and Darwen (1994) correctly pointed out, the Axelrod-style representation scheme is not the most efficient scheme. The reason why we do not use the Yao-Darwen representation scheme here is that we are restricting our attention to the case of only 3 players (oligopolists).

- **Step 2:**
  Match these three populations of ψ into l pairs of players: $\{\Xi_q\}_{q=1}^{l}$, where

$$\Xi_q = \{\psi_1^q, \psi_2^q, \psi_3^q\} \tag{15}$$

- **Step 3:**
  Let Ξ_q be applied for r periods, and calculate the profits earned by each component of Ξ_q based on Equation (4).

- **Step 4:**
  At the end of a single play (r periods), the new generation Gen_i^2 ($i = 1, 2, 3$) of the population of ψ is generated by the *canonical genetic algorithms* briefly denoted by

$$Gen_i^{t+1} = f_m \circ f_c \circ f_r(Gen_i^t), \quad i = 1, 2, 3, \tag{16}$$

  where f_m, f_c and f_r denote the genetic operators *mutation*, *crossover* and *reproduction*. The selection scheme employed is *roulette-wheel selection* and the fitness function is the profit function (11). The relevant control parameters are given in Table 5.

- **Step 5:**
  Repeat Steps 2-4 until the termination criterion is satisfied. In this paper, a pre-determined number of generations evolved (T) is chosen to be the termination criterion.

We have a few remarks on Steps 3 and 4. First, we were not simultaneously evolving the population while deriving market dynamics. Hence, the time scale of the simulation (T) is not the number of generations (Gen). For CASE A with the 8-period evolution cycle ($r = 8$), we actually evolved 250 generations, while for CASE B with the 25-period evolution cycle, we only evolved 80 generations. This naturally leads to a question: *Are these numbers of generations (periods) enough?* Due to the "*hanging valley*" well noticed by Ken Binmore, we can never be sure about this. Nonetheless, as we shall see later, all our simulation results do have a convergence result with these limited numbers of periods.

Second, based on Equation (16), the genetic algorithm is employed to evolve each population *separately*, i.e., each population is evaluated by how well it performs against itself rather than other populations. By doing this, we assumed that shops can learn the experiences *only* from those shops owned by the same oligopolist. Since in practice pricing strategies are *business secrets*, they are *not observable* and hence *not imitable*. Therefore, excluding the possibility of learning from other oligopolists' strategies seems to be a suitable approximation of the real situation.

Third, given the complexity of the oligopoly game as described in Section 2, it is not entirely clear whether *other shops' experience are relevant*. In particular, the complex dynamics of the game makes each shop's own experience *unique*. Given their different market shares, payoff vectors and local competitors, one cannot but question whether different shops can be compared on a fair ground. Nevertheless, we see no effective way to take care of all these path-dependent attributes. In fact, even in real life, people frequently simplify a complex decision process. A manager may get sacked not because of her incompetence but simply because of bad luck.

Table 6: Experimental Designs

Experiment	r	# of Runs	α
A	8	10	0.2
B	25	10	0.2

Another related and more generic issue is: *are the survivors the fittest?* As with all path-dependent dynamic systems, this is a very complicated difficult issue. Economic statistics usually show that income or wealth dynamics are path dependent, and sometimes "social justice" is required as an external force to slow down the self-reinforcing mechanism of the gap between the poor and the rich. In our usage of GA, we had a similar problem. Even though all the shops started with the same market share, the initialization process may quickly drive them apart. Since the fitness, profits, is calculated based on the market share (Equation 4), the large shops are loud in everything, which makes small shops' voices difficult to hear. So, when a small shop wins a battle with her local competitors, no one will not take notice of it because it is "puny".

4 Experimental Designs and Results

For all the experiments conducted in this study, P_h was set at "1.4", P_l "1.2" and C "1". Other control parameters of GAs were set according to Tables 5 and 6. For each set of parameters, we conducted ten independent runs, with 2000 periods for each.

4.1 Phenotypes

In the following, we shall present our simulation results in terms of the *phenotype* and the *genotype*. Before discussing the results of phenotypes, we need to clarify a few more notations. Let "W" refer to the state "*price war*" (0,0,0), "C" the state "collusive price" (1,1,1), "w" the states which are closer to "W" and "c" the states closer to "C", where "closer" is defined in terms of *Hamming distance*. Thus, "w" includes states (0,0,1), (0,1,0) and (1,0,0), and "c" includes (1,1,0), (1,0,1), (0,1,1). Since there are 30 pairs of oligopolists in each market day, a *histogram* may make the presentation easier. To do so, let p^t_W, p^t_w, p^t_c, and p^t_C denote the percentage of the pairs who, in period t, are in the states labeled with "W", "w","c", and "C" respectively. Figures 1.1-1.10 and 2.1-2.10 display the time series plot of the histogram of S_t. To see what these results suggest, a series of issues are proposed as follows.

- Is the market dynamics likely to converge?
- Is the market dynamics likely to converge to the state of collusive pricing, i.e., the state $S = (1, 1, 1)$?
- Is the market dynamics likely to converge to the state of price wars, i.e., the state $S = (0, 0, 0)$?

- Is the market dynamics likely to converge to any other states?
- Would all three firms survive to the end?

Is the market dynamics likely to converge to the state of collusive pricing? This is one of the most absorbing issues because, by Yao and Darwen (1994), "to cooperate" rather than "to defect" seems to be the most likely result in a 3-person IPD game. However, interestingly enough, none out of our 20 runs shows a convergence to the state of collusive pricing. This failure to see any result of collusive pricing is somewhat striking. It immediately drives us to the following related issue: *Is the market dynamics likely to converge to the state of price wars?* The answer seems to be *yes*. Out of the 20 runs, there are six cases whose market dynamics coverge to the state of price wars (Cases A1, B2, B6, B7, B8 and B9). Moreover, price wars seem to be more likely to occur in the case with a longer evolution cycle (Case B, r=25).

While the two questions raised above are fundamental, the market dynamics can be much richer than that. First, the market dynamics may not converge at all. Second, there is no guarantee that all three oligopolists can survive to the end. Let us start from the second point. One of the interesting features which distinguish the oligopoly game from the IPD game is that players in the oligopoly game may go extinct. This is because the existence of a player (a firm) is directly represented by her *market share*. If her market share becomes zero or infinitesimal, then the firm is effectively dead. So, *would all firms survive to the end in our simulation?* The answer is *no*. Out of the 20 cases, there are six (Cases A2, A8, A9, A10, B3, and B5) where only one firm (one oligopolist) survived to the end. In four cases (A6, B1, B4 and B10) one firm was out of the game. These ten cases shared a common pattern, i.e., *the surviving firms all charged the low price*, a phenomenon typically known as *predatory pricing* in industrial economics.

What has been simulated here is *a transition from an oligopoly industry to a monopoly or duopoly industry*. The transition process may be described as follows. Some firms initiated predatory pricing at the early stage of the game and drove their competitors out of the game. They continuously kept the price low to prevent the defeated from coming back. This behaviour is the so called "*market share first*" strategy. Genetic algorithms might help defeated firms to react to this strategy, but unlike the IPD games, the path-dependence of the game may not give players a second chance to regain back their lost market.

Now, back to the convergence issue. *Is the market dynamics likely to converge?* This issue is important because it concerns whether the parameter "number of generations" was appropriately set. If a case has a potential convergence, but we do not give them enough time to run, then the analysis based on a transient state may be misleading. By taking a quick look at the figures of the time series plots (1.1-1.10, 2.1-2.10), one may find that fluctuation appears in many plots (A2, A6, A7, A10, and B5). But, a careful examination reveals those many of these cases ended up with either a monopolist or duopolists. In this case, the action taken by the extinct firms is no longer effective. So, if we exclude the action of the nonactive firm(s) and reduce the dimension of the state to one or two, then we shall see that the market dynamics of Cases A2, A10 and B5, in effect, converges to a state of low price.

Cases A6 and A7 also converged, though, instead of a *fix point*, they converged to a *periodic cycle*. Case A6 provides us with another interesting observation. One firm went extinct in this

Table 7: Simulation Results: Phenontypes

Case	Converge?	State	Case	Converge?	State
A1	yes	price wars	B1	yes	duopolist (low price)
A2	yes	monopolist (low price)	B2	yes	price wars
A3	yes	weak price wars	B3	yes	monopolist (low price)
A4	yes	weak collusion	B4	yes	duopolist (low price)
A5	yes	weak collusion	B5	yes	monopolist (low price)
A6	yes	duopolist (periodic cycle)	B6	yes	price wars
A7	yes	periodic cycle	B7	yes	price wars
A8	yes	monopolist (low price)	B8	yes	price wars
A9	yes	monopolist (low price)	B9	yes	price wars
A10	yes	monopolist (low price)	B10	yes	duopolist (low price)

simulation, and the other two surviving firms *synchronized their pricing*, i.e., they simultaneously charged a high price followed by a low price. By doing this, firms will not invade each other's market, which is tantamount to a *tacit nonagression agreement* frequently observed in the real world. Case A7 is probably the most complicated steady state to which the market dynamics converged. It converged to a periodic cycle with periods 5. This case with the three other convergent cases, Cases A3, A4 and A5, will be discussed in detail in the next subsection. At this moment, it is enough to point out that Case A3 converged to the state "weak price wars", and Cases A4 and A5 converged to the state "weak collusion".

Table 7 summarizes what we found from these 20 simulations. None of them failed to converge. However, the results are quite diverse. From fixed points to periodic cycles, there are 7 different kinds of steady state. No state of collusive pricing come up, however, as the final outcome. Instead of cooperative behaviour, the results are overwhelmingly biased toward *predatory behaviour*. This is particularly true of Case B, where r was set at 25. The emergence and prevalence of predatory behaviour may be caused by the path dependence of the game; in particular, *the market once lost may never been regained*. This cruel fact allows firms little time to figure out the value of cooperation before taking their last breath.[5]

4.2 Genotypes

The purpose of this section is to see what kinds of pricing strategies (genotypes) acquired by firms lead to the prevalence of predatory behaviour (phenotypes), and to see how these strategies can be compared with those stereotypes of the IPD game. While dealing with the dynamics of the population of chromosomes can be a very demanding job, fortunately, in all

[5] Here, we see the interesting feature to distinguish the oligopoly game from the N-person IPD game. In the N-person IPD game, if players fail to appreciate the value of cooperation and underestimate the consequence of defection, there will always be enough time for them to learn because the payoff matrix is time invariant. In the oligopoly game, however, market shares and payoffs are time variant, and if players do not do it right in their first try, thing can be quite difficult for them. We shall see more on this in the next subsection.

Table 8: Simulation Results of Genotypes: CASE A

Case	ψ_1	ψ_2	ψ_3
A1	01101000 (0.33)	01110100 (0.24)	00000100 (0.43)
A2	01000001 (0)	00100001 (1.00)	10000011 (0)
A3	01000000 (0.36)	10011000 (0.31)	00100000 (0.33)
A4	00011000 (0.39)	11001000 (0.22)	01100011 (0.39)
A5	10000010 (0.59)	00100100 (0.32)	10100010 (0.09)
A6	00100100 (0.78)	10101000 (0)	00101000 (0.22)
A7	01001010 (0.20)	00000100 (0.59)	10011011 (0.21)
A8	00000001 (1.00)	01001001 (0)	10110110 (0)
A9	00100000 (0)	10011100 (0)	00000010 (1.00)
A10	10011000 (0)	11001010 (0)	00010000 (1.00)

Inside the bracket is the average market share of the 30 shops owned by each oligopolist. If the game converges to a fixed point in action space, the average is taken by using the market-share data from the last period. If the game converges to a periodic cycle, then the average is taken by using the data from periods of the last cycle. Notice that our computer printouts of these numbers are accurate up to the ninth decimals. However, here, we only keep the first two decimals. Therefore, we use "0.00" when the number is greater than 0.000000000, and "0" when the number is less.

our simulations, the entire population converges, i.e., the *string bias* of the population came to a value very close to 100%.[6] Or, *roughly speaking*,[7]

$$\lim_{t \to \infty} \psi_i^{1,t} = ... \lim_{t \to \infty} \psi_i^{30,t} = \psi_i, \quad i = 1, 2, 3. \tag{17}$$

Therefore, we can simply focus on the population of the last generations, $\{\psi_1, \psi_2, \psi_3\}$. Tables 8 and 9 exhibit the string to which the population converged. Based on these two tables, we shall adress the following two questions.

- What do these ψs say?
- How do these sets (ψ_1, ψ_2, ψ_3) behave?

The first question is to understand the contents of ψ from an *individual* viewpoint. However, in the context of the game, strategies cannot be well understood without taking mutual interactions into account. Therefore, the second question is posed from a *social* viewpoint.

[6] *String bias* is a measure of agreement among population. To calculate string bias, we first check the spilt between "0" and "1" at each bit position, called the *bit bias*. If the split is p%-$(1-p)$%, then the bit bias is the either p% or $(1-p)$%, depending on which one is larger. Bit bias assumes a value between 50 and 100 percent. The string bias is the average of all bit bias values. For example, in our case, it is the average of the 8 bit bias values.

[7] Given the effect of mutation disturbance, it is difficult for the population to converge completely.

Table 9: Simulation Results of Genotypes: CASE B

Case	ψ_1	ψ_2	ψ_3
B1	00001010 (0.05)	00110001 (0.95)	00000001 (0)
B2	01000001 (0.45)	00001011 (0.00)	00000001 (0.55)
B3	01100000 (0)	11011000 (0)	00000010 (1.00)
B4	10000001 (0)	00010010 (0.05)	00010000 (0.95)
B5	10010010 (0)	00000010 (1.00)	00101100 (0)
B6	00010001 (0.42)	00000110 (0.10)	00010100 (0.48)
B7	00001000 (0.55)	00000100 (0.18)	00100011 (0.27)
B8	00101001 (0.11)	00000010 (0.70)	00000001 (0.19)
B9	00000111 (0.19)	01011000 (0.00)	00000101 (0.81)
B10	11001000 (0)	01000110 (0.01)	00000000 (0.99)

The instruction to read this table is the same as Table 8.

4.2.1 Analysis from an Individual Viewpoint

What do these ψs say? At first glance, it seems difficult to discern any pattern from these two tables. The genotypes to which the three populations converged are different from one run to another. In fact, sorting through the strings shows that there are totally 25 different strategies in Table 8 and 26 in Table 9. Five strategies were used twice in Case A and 2 strategies were adopted three times in Case B. From such a low frequency of reoccurrence, one may expect more new strategies to come up in a few more runs of simulation. Since there are too many strategies shown in these two tables, it would be useful to give some general descriptions of them instead of going into the detail of each of them.

To have general description of them, we pose a series of questions for each oligopolist. The oligopolist being questioned is called the *host firm* in the question. Also, we use informal words in the questions. For example, we use "nice" in place of "to charge the high price" (cooperate), and "mean" in place of "to charge the low price" (defect).

1. If the host firm is nice to other firms, and other firms are mean to the host firm, would the host firm take revenge in the next period? I.e., to check:

$$\overbrace{\underbrace{000}\underbrace{001}\underbrace{010}\underbrace{011}\underbrace{100}_{0?}\underbrace{101}\underbrace{110}\underbrace{111}}^{Firm1},\overbrace{\underbrace{000}\underbrace{001}\underbrace{010}_{0?}\underbrace{011}\underbrace{100}\underbrace{101}\underbrace{110}\underbrace{111}}^{Firm2},\overbrace{\underbrace{000}\underbrace{001}_{0?}\underbrace{010}\underbrace{011}\underbrace{100}\underbrace{101}\underbrace{110}\underbrace{111}}^{Firm3}.$$

2. If all three firms are mean to each other, would any one like to keep it in the same way in the next period? I.e., to check:

$$\overbrace{\underbrace{000}_{0?}\,\underbrace{001}\,\underbrace{010}\,\underbrace{011}\,\underbrace{100}\,\underbrace{101}\,\underbrace{110}\,\underbrace{111}}^{Firm1,2,3}.$$

3. If all the firms are nice to each other in this period, would any one like to be mean to the others in the next period? I.e., to check:

$$\overbrace{\underbrace{000}\,\underbrace{001}\,\underbrace{010}\,\underbrace{011}\,\underbrace{100}\,\underbrace{101}\,\underbrace{110}\,\underbrace{111}_{0?}}^{Firm1,2,3}$$

4. If other firms are nice to the host firm, and the host firm is mean to them, would the host firm continuously take advantage of the others in the next period? I. e., to check:

$$\overbrace{\underbrace{000}\,\underbrace{001}\,\underbrace{010}\,\underbrace{011}_{0?}\,\underbrace{100}\,\underbrace{101}\,\underbrace{110}\,\underbrace{111}}^{Firm1},\overbrace{\underbrace{000}\,\underbrace{001}\,\underbrace{010}\,\underbrace{011}\,\underbrace{100}\,\underbrace{101}_{0?}\,\underbrace{110}\,\underbrace{111}}^{Firm2},\overbrace{\underbrace{000}\,\underbrace{001}\,\underbrace{010}\,\underbrace{011}\,\underbrace{100}\,\underbrace{101}\,\underbrace{110}_{0?}\,\underbrace{111}}^{Firm3}.$$

For all these questions "0" means a positive answer. The well-known *tit-for-tat* strategy has a sequence of answers 0-0-1-1. Using *tit-for-tat* as a benchmark, a strategy with any "1" appearing in the first two positions can be considered *merciful* or *gentle*; a "0" appearing in the *last two posotions* can be considered *aggressive*. Since each strategy may have different answers, for each questions we simply calculated the proportion of the firms who have a positive answer. It was found that the statistics for Case A are 0.73, 0.63, 0.77, 0.60, and for Case B are 0.83, 0.87, 0.60, 0.76. While some of these statistics are not statistically significantly different from 0.5 at 0.05 significance level, all of them are greater than 0.5. Therefore, loosely speaking, strategies evolved from our two sets of oligopoloy games are not *gentle*, but *a little aggressive*. This result is not surprising and is consistent with the prevalence of predatory behaviour observed.

Of course, not all strategies can help firms to survive well. To have a vivid picture of this, we take an average of the market shares of 30 shops own by each oligopolist. If the game converges to a fixed point in action space, this average is taken by using the market-share data from the last period. If the game converges to a periodic cycle, then the average is taken by using the data from periods of the last cycle. These average market shares are reported along with strategies in Tables 8 and 9. From 0 to 1, these two tables show a great dispersion of market shares among different runs. Notice that the initial market shares for each firm is *one third*, but only a few runs, e.g., Case A3, ended up with firms with equal market shares. Instead, the normal ecology seems to be the one with unbalanced market power.

Table 10: Simulation Results: Attractors, CASE A

Case	#	Attractors
A1	2	000, 100
A2	3	$(000 \rightarrow 001 \rightarrow 100 \hookleftarrow)$, 010, 111
A3	1	$(100 \rightarrow 010 \rightarrow 001 \hookleftarrow)$
A4	1	$(001 \rightarrow 011 \rightarrow 100 \rightarrow 110 \hookleftarrow)$
A5	1	$(000 \rightarrow 101 \rightarrow 010 \rightarrow 011 \hookleftarrow)$
A6	1	$(000 \rightarrow 010 \rightarrow 111 \hookleftarrow)$
A7	1	$(000 \rightarrow 001 \rightarrow 100 \rightarrow 101 \rightarrow 010 \hookleftarrow)$
A8	1	$(010 \rightarrow 001 \hookleftarrow)$
A9	1	$(010 \rightarrow 100 \hookleftarrow)$
A10	1	$(000 \rightarrow 110 \rightarrow 010 \hookleftarrow)$

The sign "$\hookleftarrow$" refers to the start of another cycle.

At this point, one may be tempted to ask *what the most competitive strategy is.* However, in an *coevolutionary* context, this question is not really well-defined, since everything depends on everything else. As a result, even though the strategy is completely the same, depending on the strategies it competes against, the associated payoff can be quite different. For example, both the third oligopolists in Cases B1 and B8 used the strategy "00000001". Nevertheless, the former went extinct, while the latter had a 20% market share.

Although the best strategy may be ill-defined, there is a noticeable difference in the characteristic of two extremes of the strategies, i.e., the ones leading to a great success (a 100% market share), and the ones leading to a great failure (a 0% market share). The difference lies in the frequency of the bit "1" appearing in the strategy. The average number of "1" appearing in the former cases is 1.16, while it is 2.88 for the latter. More generally, we run a *regression* of the average market share against the frequency of "1" appearing in each cases, and the results are as follows.

$$Market - Share = \begin{cases} 0.72 - 0.15(\#\ of\ 1), & \text{for } CASE\ \ A, \\ 0.78 - 0.22(\#\ of\ 1), & \text{for } CASE\ \ B. \end{cases} \tag{18}$$

In both cases, the regression coefficient is negative. Therefore, the regression results does suggest that "cooperate" is a risky play.

4.2.2 Analysis from a Social Viewpoint

How do the sets (ψ_1, ψ_2, ψ_3) *behave?* When three strategies are grouped together, given an initial condition, they will generate a sequence of actions $S_t(= (a_1^t, a_2^t, a_3^t))$, and $a^t \in \{0, 1\}$. Since all cases converged, we shall only focus on the asymptotic behaviour of S_t given a specific set (ψ_1, ψ_2, ψ_3), i.e., the *attractors* of S_t. Tables 10 and 11 displays these attractors. For the case who does not have a unique attractor, we list all of them, while by putting the one which was actually visited in the first place. For example, Case A1 has two attractors "000" and

Table 11: Simulation Results: Attractors, CASE B

Case	#	Attractors
B1	3	000, 010, 100
B2	2	000, 111
B3	2	(010 $\rightarrow$ 100 $\hookleftarrow$), (001 $\rightarrow$ 110 $\hookleftarrow$)
B4	2	(000 $\rightarrow$ 100 $\hookleftarrow$), 011
B5	2	(000 $\rightarrow$ 100 $\rightarrow$ 001 $\hookleftarrow$), 011
B6	2	000, (011 $\rightarrow$ 101 $\hookleftarrow$), 011
B7	2	000, 100
B8	2	000, 100
B9	2	000, 101
B10	1	100

The sign "$\hookleftarrow$" refers to the start of another cycle.

"100", but only "000" was realized as the final state of the game.

There are two kinds of attractors appearing. The first kind of attractors is *fixed points*. It is interesting to note that the *only* fix point which was realized as the final state of the outcome is "000", i.e., *price wars* (Cases A1, B1, B2, B6, B7, B8, and B9). This result is quite obvious: if one firm decide to charge the low price from now on, then the only possibility for other firms to survive is to simply join in. Consequently, *if "111" does not appear, "000" can be the only fixed point.* However, there is one exception, i.e., Case B10, whose final state is "100". But, as one would expect, the first firm did not survive to the end, and hence the first bit of this state is clearly superfluous.

While the results suggest that it was unlikely for all the firms to charge the high price all the time ("111"), it does not necessarily mean that they would always charge the low price ("000"). The second kind of attractors, i.e., periodic cycles with different periods, indicates something in between. In these attractors, firms could charge the high price at the same time, or at different times. Whichever the case, *no firm would charge the low price at all times.*

Let us consider the first kind of possibility, i.e., firms would charge the high price at the same time. Case A6 is the only example. In this case, the cycle go through the following three states: "000", "010" and "111". Nevertheless, since the second firm did not survive to the end, the cycle can be reduced to a one in the two-dimension state space: "00", "00" and "11". In such a cycle, the two surviving firms would concurrently charge the low price for two periods, then switch to the high price in the following period, and switch back the low price again. This behaviour can be identified as the *nonagression agreement* as discussed in Section 2. In Case 6, this nonagression agreement was an emergent behaviour, i.e., the agreement was not made at the begining of the game. It was achieved in a complex path-dependent dynamics. When the achievement was made, the first firm was already a big firm (a 78% market share), and the third firm turned out to be a medium one (a 22% market share).

For the second kind of possibility, consider Case A4. The attractor of Case A4 has a periodic cycle with period four. This cycle go through the following four states: "001", "011", "100", and "110". In such a cycle, the sequences of actions taken by the three firms are "0-0-1-1", "0-1-0-1" and "1-1-0-0" respectively. In other words, in a four-period cycle, all the firms charged both the high price and the low price twice, while at different times.

As we briefly mentioned above, "charging high prices *in turn*" in practice can be considered a collectively intelligent behaviour, because it is a *smart* idea to avoid the charge of the anti-trust law. In this coordination, on the table firms make a *fictitious competition*, while under the table they reach a tacit agreement that one half of the time they can charge a high price. By the tacit agreement, they are guaranteed to not lose market to each other. Hence, this kind of behaviour can also be classified as a type of *nonagression agreement*. Moreover, since they are charging high prices at different times, unlike Case A6, it is very difficult to verify their anti-competition behaviour. Therefore, "*charging the high price in turn*" is a even more sophisticated design of cooperation. In addition to Case A4, this type of emergent behaviour is also observed in Cases A3, A5, and A7. Therefore, *cooperative behaviour does emerge in our simulations, while it is manifested in a much more subtle way as opposed to the IPD game.*[8]

The emergence of periodic cycles can also help us in accounting for the three stylized phenomena of the oligopolistic industry, as we summarized in the first section. The periodic cycle of Case A6,

$$\underbrace{000}_{Price\ \ War} \rightarrow 010 \rightarrow \underbrace{111}_{Collusive\ \ Pricing} \hookleftarrow, \tag{19}$$

is somewhat close the first stylized fact, i.e., collusive pricing is frequently interrupted by the occurrence of predatory pricing. The periodic cycle of Case A4,

$$001 \rightarrow 011 \rightarrow 100 \rightarrow 110 \hookleftarrow, \tag{20}$$

is pretty much about the second stylized fact, i.e., a dispersion of prices can persistently exist. As to the third stylized fact, i.e., firms continuously switch between the high price and the low price, it is a common property of the second kind of attractors. Of course, the real pricing patterns of oligopoly is far more complex than what we have shown here. For example, in the real world, pricing series may not have any *regular cycle* at all. However, given the simplicity of our model, our results are encouraging enough to suggest that this is a good *starting point* to advance the study of oligopolists' behaviour.

5 Concluding Remarks

In this study, the genetic algorithm was applied to an oligopoly game. While like in the well-known IPD game firms (players) encounter a similar subtle decision as *to defect* or *to cooperate*, the Markov-process characterization of firms' market shares makes the oligopoly game a non-trivial generalization of the IPD game. First, depending on the initial market shares and transition rules, the payoff matrix is time-variant and state-dependent. Second, as a result, the inequalities which define the IPD game may be violated as time goes on, which means even though an oligopoly game can satisfy the conditions of the IPD game, the learning process

[8]Apart from not running against the anti-trust law, firms can benefit from this tacit agreement if both the sixth and the seventh condition of the IPD game is violated, as been discussed in Section 3.

randomly initiated may fail them in a later stage. Due to these extensions, one may hence conjecture that results of oligopoly games can be too rich to be predicted by a single factor, such as the number of players, and this conjecture is confirmed by this paper.

In our simulations, we show that, even in the *three-player case*, collusive pricing (cooperative behaviour) is not the dominating result as one may expect from the standard 3-person IPD game. The results of our 20 simulations are quite divergent, but together they give quite a vivid replication of what one may observe from a real oligopoly industry. The three characteristics of the pricing patterns are captured by many of our simulations (Cases A3, A4, A5, and A7). In addition to this, in many of our simulations, we also experience a transition from an oligopoly industry to a monopoly (A2, A8, A9, A10, B3, B5) or duopoly industry (A6, B1, B4, B10). One thing failed to be generated in our simulation is persistent collusive pricing. Persistent pricing wars instead seem to the dominating outcomes. While a more advanced model could be attempted with complicated GAs, we very much doubt if the rich nature of the oligopoly industry would change due to these sophistication.

References

[1] Abreu, D., D. Pearce and E. Stacchetti (1986), "Optimal Cartel Equilibria with Imperfect Monitoring," *Journal of Economic Theory*, 39, pp.251-269.

[2] Beaufils, B. J.-P. Delahaye and P. Mathieu (1998), "Complete Classes of Strategies for the Classical Iterated Prisoner's Dilemma," in V. W. Porto, N. Saravanan, D. Waggen and A. E. Eiben (eds.), *Evolutionary Programming VII*, pp. 32-41.

[3] Fudenberg, D. and J. Tirole (1989), "Game Theory for Industrial Organization: Introduction and Overview," in R. Schmalensee, and R. D. Willig (eds.), *Handbook on Industrial Organization*, North Holland.

[4] Green, E. J. and R. H. Porter (1984), "Non-cooperative Collusion under Imperfect Price Information," *Econometrica*, 52, pp.87-100.

[5] Midgley, D. F., R. E. Marks, and L. G. Cooper (1997), "Breeding Competitive Strategies," *Management Sciences*, Vol. 43, No. 3, pp. 257-275.

[6] Shapiro, C. (1989), "Theories of Oligopoly Behavior," in R. Schmalensee, and R. D. Willig (eds.), *Handbook on Industrial Organization*, North Holland.

[7] Scherer, F. M. and D. Ross (1990), *Industrial Market Structure and Economic Performance*, 3rd edition, Houghton Mifflin.

[8] Yao, X. and P. J. Darwen (1994), "An Experimental Study of N-Person Iterated Prisoner's Dilemma Games," *Inoformatica*, Vol. 18, pp. 435-450.

case B. alpha=0.2, beta=0

mutation rate = 0.001,r=8,Ph=1.4

Figure A1. The Distribution of States (seed=78273.38)

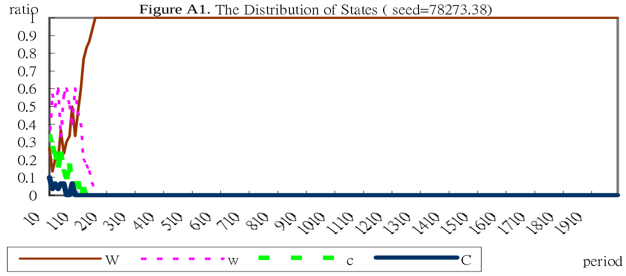

Figure A2. The Distribution of States (seed=79297.74)

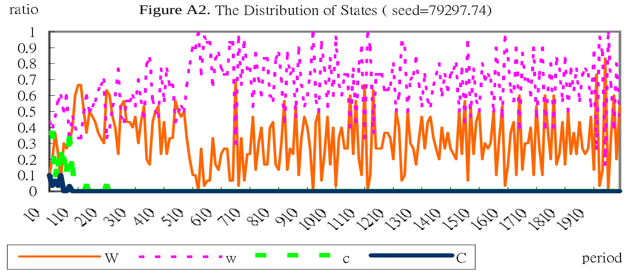

Figure A3. The Distribution of States (seed=79638.12)

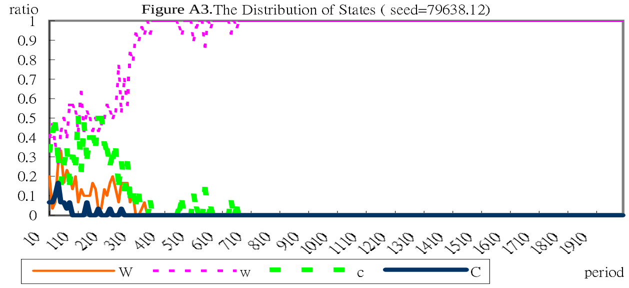

Figure A4. The Distribution of States (seed=80151.78)

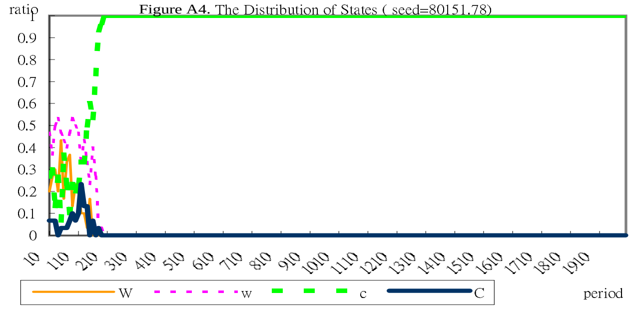

Figure A5. The Distribution of States (seed=80466.94)

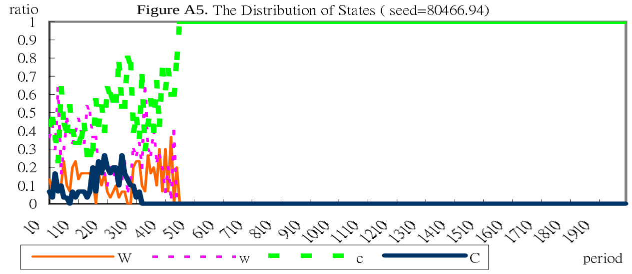

Figure A6. The Distribution of States (seed=57702.8)

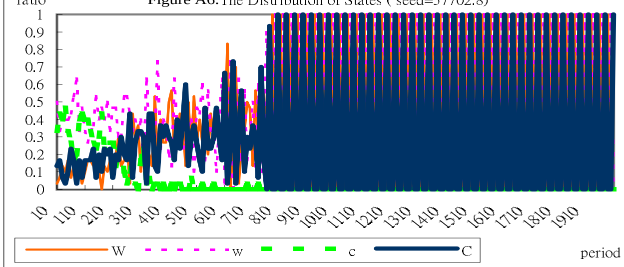

Figure A7. The Distribution of States (seed=57869.27)

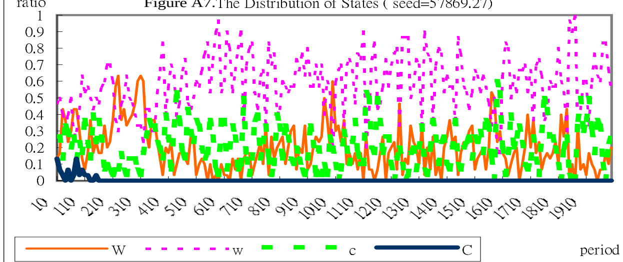

Figure A8. The Distribution of States (seed=58241.01)

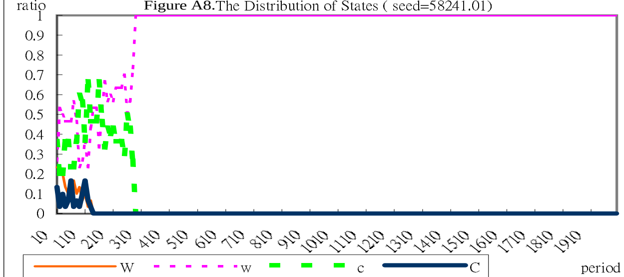

Figure A9. The Distribution of States (seed=58507.18)

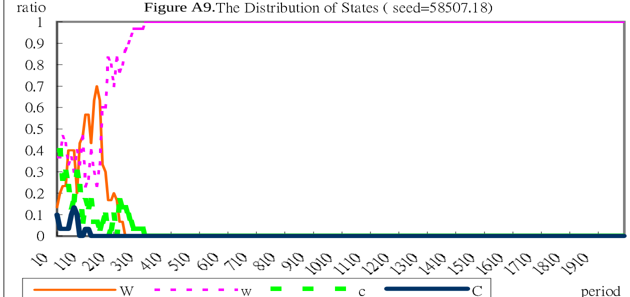

Figure A10. The Distribution of States (seed=58746.71)

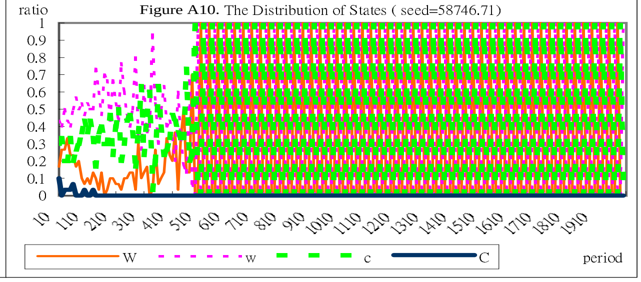

1011-1.xls

case B. alpha=0.2, beta=0

mutation rate = 0.001,r=25,Ph=1.4

Figure B1. The Distribution of States (seed=78273.38)

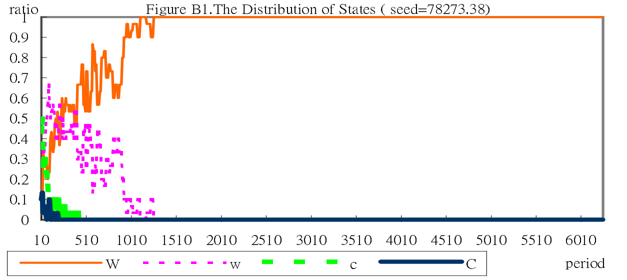

Figure B2. The Distribution of States (seed=79297.74)

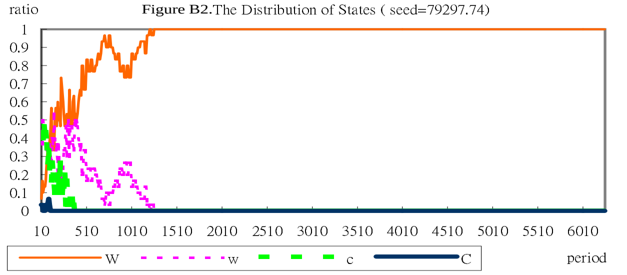

Figure B3. The Distribution of States (seed=79638.12)

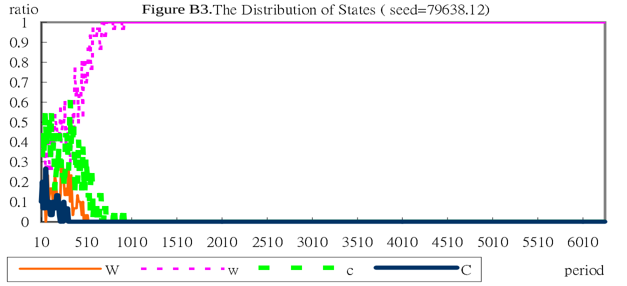

Figure B4. The Distribution of States (seed=80151.78)

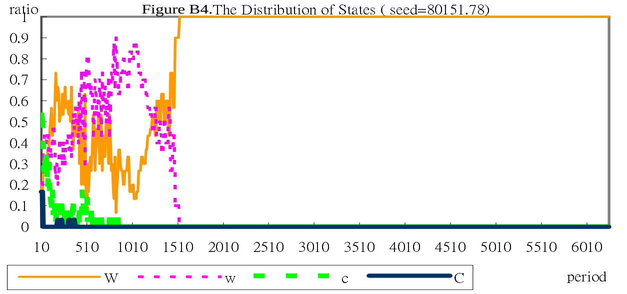

Figure B5. The Distribution of States (seed=80466.94)

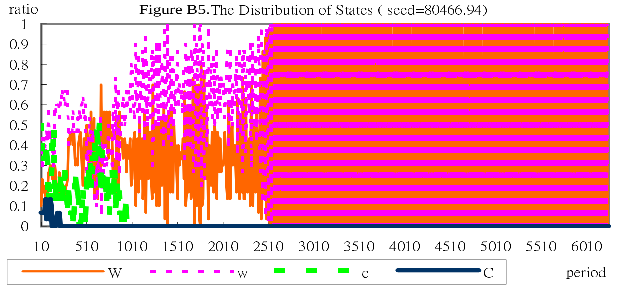

Figure B6. The Distribution of States (seed=57702.8)

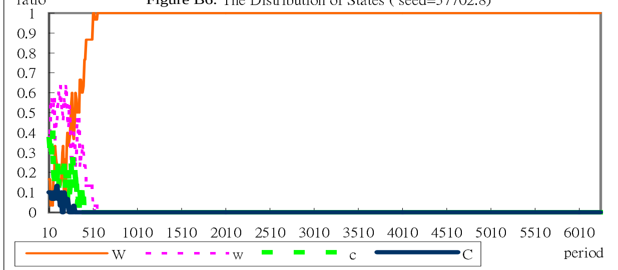

Figure B7. The Distribution of States (seed=57869.27)

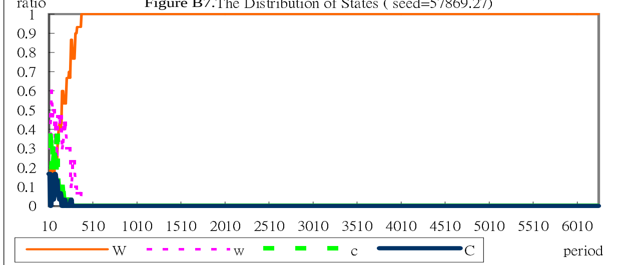

Figure B8. The Distribution of States (seed=58241.01)

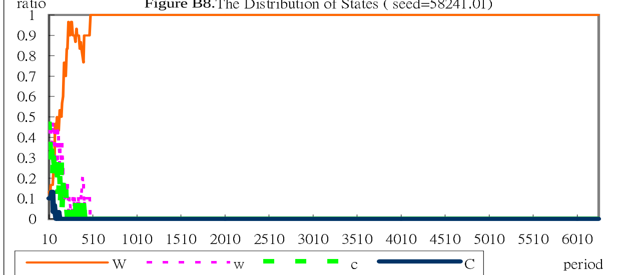

Figure B9. The Distribution of States (seed=58507.18)

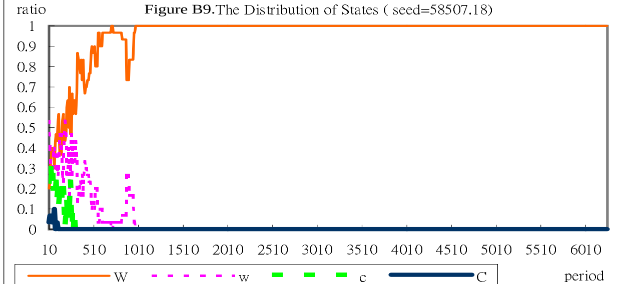

Figure B10. The Distribution of States (seed=58746.71)

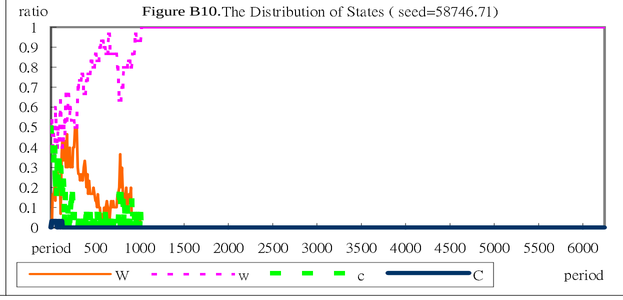

2011-1.xls