

Simulating the Ecology of Oligopolistic Competition with Genetic Algorithms

Shu-Heng Chen and Chih-Chi Ni

National Chengchi University, Taipei, Taiwan

Abstract. In economics, the n -person oligopoly game and the n -person IPD game are often considered close in spirit. Our analytical framework shows that this is not the case owing to the path dependence of the pay-off matrix of the oligopoly game. By simulating the evolution of a three-person oligopoly game with genetic algorithms, we explore the significance of the path dependence property to the rich ecology of oligopoly. The emergent behavior of oligopolists in the simulations indicates how the path dependence nature may shed light on the phenotypes and genotypes coming into existence. The features shown in this research can be further exploited in more practical contexts so that nontrivial policy issues in industrial economics can be seriously tackled.

Keywords: Coevolution; Genetic algorithms; Market-share-first strategy; Oligopoly; Predatory pricing; State-dependent Markov chain

1. Motivation and Introduction

In industrial economics, modeling a market consisting of only a few firms, i.e., oligopolistic industry, is a tricky task. 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 in textbooks as an effective abstraction of the oligopoly pricing problem. 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

Received 12 February 1999

Revised 9 December 1999

Accepted 3 April 2000

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

players in the n -person oligopoly game are facing a pretty similar situation to those in the n -person IPD 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 three, it is very likely to see the emergence of collusive pricing (cooperation)*. Real data, however, shows that even in a three-oligopolist industry the observed pricing pattern is not that simple (Midgley et al, 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, while 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 behavior, gravitate toward price warfare. (Scherer and Ross, 1990, p. 199; italics added)

Nonetheless, ‘*virtually anything can happen*’ is not the property which one may experience from a three-person IPD game (see Yao and Darwen, 1994, Fig. 5), and this raises two questions:

- 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 *twofold*. First, contrary to some people’s presumption, we show that in general the n -person oligopoly game is *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 three oligopolists. In this game, the pay-off 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 although in general these two games are not close in spirit, *there does exist a trivial path* on which the two games are *effectively* the same. Therefore, we

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 Midgley et al (1997).

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. 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 et al, 1986), as many conventional studies did. Instead, we are asking: *other things being equal, can we still have a rich ecology of oligopoly?* We consider such an effort a search for a more *fundamental cause*, and this constitutes the second contribution of the paper.

In Section 3, we illustrate the use of genetic algorithms (GAs) to model the adaptive behavior of oligopolists. The application of GAs to the oligopoly game is nothing new.³ Midgely et al (1997) pioneered this line of research. While we follow the ideas employed in their paper in many aspects, there is an important distinction between the two studies. 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 have already done an excellent job. What concerns us instead is to use GAs to allow the oligopoly game to write history on its own and to see how rich the ecology can be. Therefore, it is sufficient to just let our GA-based oligopolists learn from their own experience and acquire expertise without being exposed to real data.

As we will see in Section 4, what was derived 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 oligopoly game. Some interesting patterns such as the *market-share first* strategy, *nonaggression agreement*, *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 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* :

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

³ For an extensive survey on the application of GAs to game theory, the interested reader is referred to Marks (2000).

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 the following equation:

$$N_{t+1} = N_t M_t \quad (2)$$

To see the effect of price competition on the market share dynamics, the transition probability m_{ij}^t is assumed to be dependent on the pricing strategy vector S_t . If the three firms charge the same price, then M_t is an *identity matrix*. Furthermore, if firm i charges P_h , then it will lose $\frac{1}{2}\alpha \times 100\%$ of its consumers to each of the two firms j and k , who charge P_l . Furthermore, if firms i and j charge P_h , then each of them will lose $\alpha \times 100\%$ of their consumers to firm k , who charges P_l . These assumptions can be summarized by the following transition matrices:

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

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

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

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 the number of customers $n_i = \sum_{i=1}^3 n_i^t$.

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):

$$\pi_i^s = (P_i^s - C)n_i^s \quad (4)$$

where P_i^s is the price charged by firm i at period s , n_i^s the number of customers, and C the 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 pay-off matrix used to define an n -person IPD game (Yao and Darwen, 1994). However, due to the market share dynamics, the pay-off 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 to ‘cooperate’ (C) means to ‘charge the high price for all r periods’ and to ‘defect’ (D) means to ‘charge the low price for all r periods’.

We can now work out the first-round *pay-off matrix* employed by Yao and Darwen (1994). In our case (three oligopolists), there are six elements in the pay-off matrix, namely C_i and D_i ($i = 0, 1, 2$). Here, C_i (D_i) denotes the pay-off 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:

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

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

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

and

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

where

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

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

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:

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

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

Table 1. Parameters and pay-offs: the first round of the game

Set	P_H	P_L	C	α	r	D_2	D_1
1	1.4	1.2	1	0.2	8	3.47	2.07
2	1.4	1.2	1	0.2	25	13.40	7.10

Set	D_0	C_2	C_1	C_0	$\frac{1}{2}(D_2 + C_1)$	$\frac{1}{2}(D_1 + C_0)$
1	1.6	3.2	1.33	1.33	2.4	1.70
2	5	10	1.60	1.60	7.50	4.35

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

Based on the derived pay-off 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, Fig. 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) $C_2 > \frac{1}{2}(D_2 + C_1)$, and (7) $C_1 > \frac{1}{2}(D_1 + C_0)$.

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 the others 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 behavior. It is interesting to note that in reality oilgopolists 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 behavior, it is useful to take notice of this emergent intelligence in our simulations which will be discussed later.

By equations (9)–(11), the pay-off 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 pay-offs 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.

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. The other signs mean the condition is weakly violated (=) or strongly violated (<).

Table 3. Parameters and pay-offs: the second round of the game

Round 1/2	D_2	D_1	D_0	C_2	C_1	C_0	$\frac{1}{2}(D_2 + C_1)$	$\frac{1}{2}(D_1 + C_0)$
<i>Parameter set 1</i>								
D_2	4.58	4.34	4.26	8.53	3.55	3.55	4.06	3.94
D_1	3.75	2.34 2.93	2.27	4.53	1.89	1.89	2.82	2.11 2.41
D_0	3.47	2.07	1.60	3.20	1.33	1.33	2.40	1.70
C_2	3.47	2.07	1.60	3.20	1.33	1.33	2.40	1.70
C_1	2.91	0.35 1.51	0.27	0.54	0.22	0.22	1.57	0.29 0.87
C_0	2.91	0.93	0.27	0.54	0.22	0.22	1.57	0.58
<i>Parameter set 2</i>								
D_2	14.99	14.97	14.96	29.92	4.77	4.77	9.88	9.87
D_1	13.80	7.50 10.64	7.49	14.98	2.39	2.39	8.10	6.51 4.94
D_0	13.41	7.10	5.00	10.00	1.59	1.59	7.50	4.35
C_2	13.41	7.10	5.00	10.00	1.59	1.59	7.50	4.34
C_1	12.61	0.03 6.31	0.02	0.04	0.01	0.01	6.31	0.02 3.16
C_0	12.61	3.17	0.02	0.04	0.01	0.01	6.31	1.59

So far, we have only worked out the pay-off vector of the first round of the game. The pay-off vector of the second round, and the rounds after, is a little intriguing. Due to the dynamics of market shares, the pay-off vector is not independent of what happened in the first round. In other words, the pay-off 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 pay-off vector too. Since there are six nonredundant histories (paths) in the first round, each with a follow-up pay-off vector, the second-round pay-offs of the game can be represented by a 6×6 pay-off table as shown in Table 3.

Table 3 exhibits the pay-off 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 pay-off matrix in three ways. First of all, in the first round of the game the pay-off is symmetric, while in the second it depends. For example, if in the first round the three firms all charge the high price, or the low

Table 4. Parameter sets and testing results

Inequality/initial condition	D_2	D_1	D_0	C_2	C_1	C_0
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. The other signs mean the condition is weakly violated (=) or strongly violated (<). Signs in brackets refer to the reversals of the pay-off inequality in the second round of the game.

price for that matter, then in the second round the initial condition for them is the same, and hence their pay-off vectors refer to the same row led by C_2 or D_0 in Table 3. In this case, the symmetry of pay-offs remains unchanged. However, if in the first round two firms charge the high price, and one firm charges the low price, then in the second round the pay-off vector to the firms who charged the high price is the one led by C_1 , while the pay-off vector to the firm who charged the low price is the one led by D_2 . According to Table 3, these 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 pay-off vectors, it is interesting to note that in some cases the pay-off is *not unique*. For example, in Table 3, if the initial state for the player is D_1 , and the current state is also D_1 , then the pay-off 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 large market share cooperates, then ‘defection’ can be more profitable because of the proportion of the market one could seize.

Last, the reversal of the pay-off inequality. Based on Table 3, we checked the pay-off 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 D_2 or D_1 , instead of ‘ $D_2 > C_2$ ’, we have ‘ $D_2 < C_2$ ’, and when the initial condition is C_1 , ‘ $C_2 > \frac{1}{2}(D_2 + C_1)$ ’. The reversal of the first condition is *critical*, because if two players choose C , then the dominant 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 behavior, i.e., ‘*false defection*’ as discussed above.

In sum, the oligopoly game is not an n -person IPD game. It is nonetheless related to the n -person IPD game in a subtle way. In particular, among the 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 three-person oligopoly game may be quite different from that of a three-person IPD game. It is therefore interesting to know whether the *path dependence property* of the oligopoly game can generate more complex patterns of emergent behavior than the n -person IPD game, and if so, *what are they* and *what do they mean*?

3. Modeling the Adaptive Behavior of Oligopolists with GAs

In this study, we simulated 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 their pricing strategies, each shop interacted locally with other shops' pricing. The pricing strategy of each shop was represented by a *binary string*, and thus the pricing strategies of all the shops owned by the same oligopolist were a *population of binary strings*. The GA 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:

$$\phi : \Omega \longrightarrow \{0, 1\} \quad (12)$$

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

$$\psi : \Omega_k \longrightarrow \{0, 1\} \quad (13)$$

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 action 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 that concerns us is the smallest value of k which can reasonably replicate the price dynamics of the oligopoly industry and, as we shall see later, setting k to equal 1 is good enough to achieve this goal.⁴ In Ω_1 , ψ can be coded with an 8-bit string. For example, an 8-bit string $b_1b_2 \dots b_8$ means that if state j ($j = 1, 2, \dots, 8$) occurs, the oligopolist will take action b_j ($b_j = 0, 1$).⁵ The eight states are ordered as in the following sequence:

$$\begin{array}{cccccccc} \underbrace{000}_1 & \underbrace{001}_2 & \underbrace{010}_3 & \underbrace{011}_4 & \underbrace{100}_5 & \underbrace{101}_6 & \underbrace{110}_7 & \underbrace{111}_8 \end{array}$$

⁴ 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.

⁵ 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 three players (oligopolists).

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

Each state is represented by a 3-bit string. From left to right, the first bit refers to the action taken by firm 1 in the previous period, the second bit refers to the action taken by firm 2, and so on. For instance, state '5' encoded as '100' means that firm 1 charged the high price (cooperated), but firms 2 and 3 charged the low price (defected) in the last period. If $b_5 = 0$ for firm 1, then firm 1 will take revenge by charging the 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 ($i = 1, 2, 3$):

$$Gen_i^1 = \{\psi_i^1, \dots, \psi_i^q, \dots, \psi_i^l\} \quad (14)$$

where l is the size of population (the number of chromosomes).

• **Step 2:**

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

$$\Xi_q = \{\psi_1^q, \psi_2^q, \psi_3^q\} \quad (15)$$

• **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

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

where f_m, f_c and f_r denote the genetic operators *mutation*, *crossover*, and *reproduction*, and f_i is the *immigration* operator.⁶ 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.

⁶ The immigration operator replaces an existing chromosome with one which is totally randomly generated. The idea of employing this operator is to enable one to observe the evolutionary stability of a particular equilibrium achieved.

- **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 the 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.⁷

Second, based on equation (16), GAs were employed to evolve each population *separately*; i.e., each population is evaluated by how well it performs against itself rather than against other populations. By doing this, we assume that shops can extract experience *only* from those owned by the same oligopolist. Since in practice pricing strategies are *business secrets*, they are *not observable* and hence *impossible to imitate*. Therefore, excluding the possibility of learning from other oligopolists’ strategies seems to be a fair 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’ experiences are relevant*. In particular, the complex dynamics of the game makes each shop’s own experience *unique*. Given their different market shares, pay-off 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.

Another related and more generic issue is: *are the survivors the fittest?* As with all path-dependent dynamic systems, this is a very complicated 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 poor and rich. In our application of GAs, 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 against her local competitors, it will probably go unnoticed because such a victory 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

⁷ The term ‘hanging valley’ is given by Ken Binmore in the keynote speech at the 1997 *International Conference on Computer Simulation and Social Sciences* held in Italy. His speech is collected in Conte et al (1997). Binmore used the term to assert that it may take a very, very long time before the state collapses or transits to other states.

Table 6. Experimental designs

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

each set of parameters, we conducted 10 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 pricing’ (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 *histrogram* may make the presentation easier. To do so, let p_W^t , p_w^t , p_c^t , and p_C^t denote the percentage of the pairs who in period t are in the states labeled with ‘W’, ‘w’, ‘c’, and ‘C’ respectively. Figures 1(a–j) and 2(a–j) 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:

- Are the market dynamics likely to converge?
- Are the market dynamics likely to converge to the state of collusive pricing, i.e., the state $S = (1, 1, 1)$?
- Are the market dynamics likely to converge to the state of a price war, i.e., the state $S = (0, 0, 0)$?
- Are the market dynamics likely to converge to any other states?
- Will all three firms survive to the end?

Are the market dynamics likely to converge to the state of collusive pricing? By Yao and Darwen (1994), ‘to cooperate’ rather than ‘to defect’ seems to be the most likely result in a three-person IPD game. Yet, none of our 20 runs shows a convergence to the state of collusive pricing. This result is somewhat striking. It immediately drives us to the following related issue: *Are the market dynamics likely to converge to the state of a price war?* The answer seems to be *yes*. Out of the 20 runs we conducted, there are six cases in which the market dynamics converge to the state of a price war (cases A1, B2, B6, B7, B8, and B9). Moreover, price wars seem to be more likely to occur in cases 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. Since in the oligopoly game the existence of a player (a




Fig. 1. For caption see facing page.

firm) is directly represented by her *market share*, if her market share comes down to zero or becomes infinitesimal, then the firm is effectively dead. So, *would all firms survive to the end in our simulations?* 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 10 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.

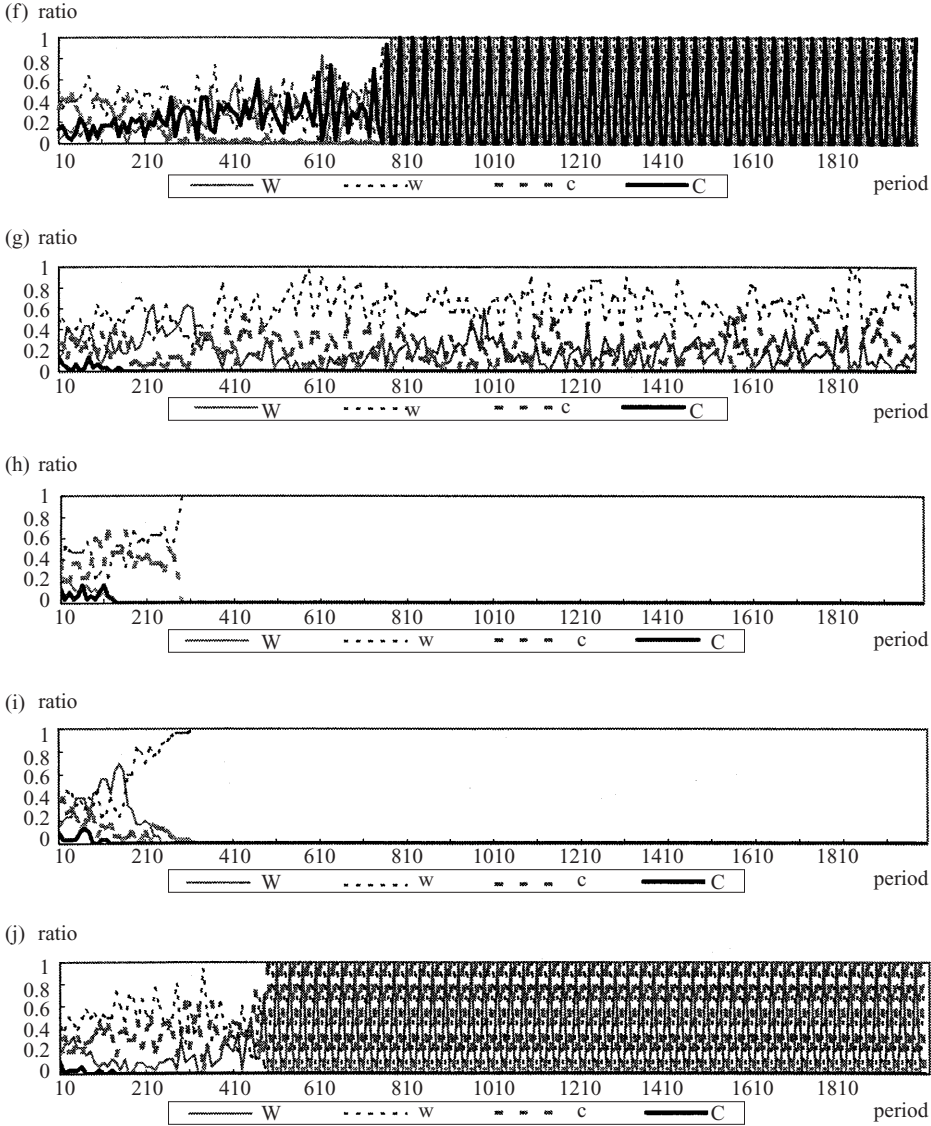


Fig. 1. The distribution of states: (a) seed = 78273.38; (b) seed = 79297.74; (c) seed = 79638.12; (d) seed = 80151.78; (e) seed = 80466.94; (f) seed = 57702.8; (g) seed = 57869.27; (h) seed = 58241.01; (i) seed = 58507.18; (j) seed = 58746.71.

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 is what we call the 'market share first' strategy. GAs might help defeated firms to react to this strategy, but unlike




Fig. 2. For caption see facing page.

the IPD games the path dependence of the oligopoly game may not give players a second chance to regain their lost market.

Now, back to the convergence issue. *Are the market dynamics likely to converge?* To answer this question, the number of generations needs to be appropriately set. If a case has a potential to converge, but we do not give it 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(a-j), 2(a-j)), one may find that fluctuation appears in many plots (A2, A6, A7, A10, and B5). But, a careful examination reveals that many of these cases ended up with either a monopolist

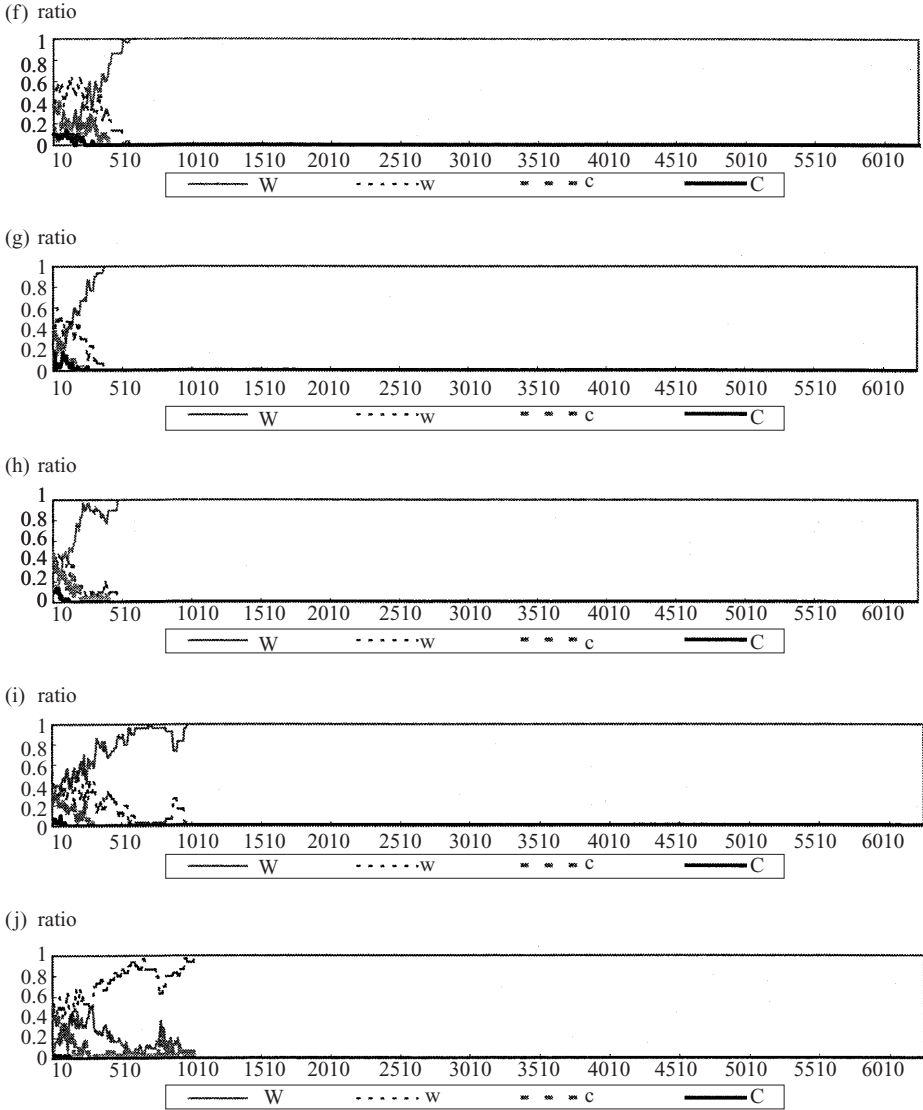


Fig. 2. The distribution of states: (a) seed = 78273.38; (b) seed = 79297.74; (c) seed = 79638.12; (d) seed = 80151.78; (e) seed = 80466.94; (f) seed = 57702.8; (g) seed = 57869.27; (h) seed = 58241.01; (i) seed = 58507.18; (j) seed = 58746.71.

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, converge to a state of the 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 simulation, and the other two surviving firms *synchronized their pricing*; i.e., they simultaneously charged the high price followed by the low

Table 7. Simulation results: phenotypes

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)

price. By doing this, firms will not invade each other’s market, which is tantamount to a *tacit nonaggression 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 five periods. This case, along with 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 fail to converge, but the results are quite diverse. From fixed points to periodic cycles, there are totally seven different kinds of steady state. No state of collusive pricing comes up as the final outcome, however. Instead of cooperative behavior, the results are overwhelmingly biased toward *predatory behavior*. This is particularly true of case B, where r was set at 25. The emergence and prevalence of predatory behavior may have been caused by the path dependence of the game. The fact that *markets once lost may never be regained* allows firms little time to realize the value of cooperation before they take their last breath.⁸

4.2. Genotypes

The purpose of this section is to see what kinds of pricing strategies (genotypes) the firms acquired led to the prevalence of predatory behavior (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, the entire population fortunately converged in all our

⁸ Here, we see the interesting feature that distinguishes 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 pay-off matrix is time invariant. In the oligopoly game, however, market shares and pay-offs are time variant, and if players do not do it right in their first try things 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 parentheses is the average market share of the 30 shops owned by each oligopolist. If the game converged to a fixed point in action space, the average was taken by using the market share data from the last period. If the game converged to a periodic cycle, then the average was 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 decimal. Here, however, we only keep the first two decimal. Therefore, we use '0.00' when the number is larger than 0.000000000, and '0' when it is smaller.

simulations; i.e., *string bias* of the population was very close to 100%.⁹ Or, *roughly speaking*,¹⁰

$$\lim_{t \rightarrow \infty} \psi_i^{1,t} = \dots \lim_{t \rightarrow \infty} \psi_i^{30,t} = \psi_i, \quad i = 1, 2, 3 \quad (17)$$

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 address the following two questions:

- What do these ψ s 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. In the context of the game, however, strategies cannot be well understood without interaction being taken into account. Therefore, the second question is posed from a *social* viewpoint.

4.2.1. Analysis from an Individual Viewpoint

What do these ψ s say? At first glance, it seems difficult to discern any pattern from these two tables. The genotypes to which the three populations converge 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 two strategies were adopted three times in case B. From such a low frequency of recurrence, one may expect more new strategies to come up in a few more runs of simulation. Since there are too many strategies

⁹ *String bias* is a measure of agreement among the population. To calculate string bias, we first check the split between '0' and '1' at each bit position, called *bit bias*. If the split is $p\%$ -($1-p$)%, then bit bias is either $p\%$ or $(1-p)\%$, depending on which one is larger. Bit bias assumes a value between 50% and 100%. String bias is the average of all bit bias values. In our case, string bias is the average of the 8-bit bias values.

¹⁰ 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)

Instructions for reading this table are the same as for Table 8.

shown in these two tables, it seems advisable to give a general description of these strategies instead of going into the detail of each of them.

To do so, we pose a series of questions for each oligopolist. The oligopolist in question is called the *host firm*. For the sake of brevity, 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 the other firms, but the other firms are mean to the host, will the host firm take revenge in the next period? That is, to check:

$$\overbrace{\underbrace{000} \underbrace{001} \underbrace{010} \underbrace{011} \underbrace{100} \underbrace{101} \underbrace{110} \underbrace{111}}^{\text{Firm1}} \\ 0?$$

$$\overbrace{\underbrace{000} \underbrace{001} \underbrace{010} \underbrace{011} \underbrace{100} \underbrace{101} \underbrace{110} \underbrace{111}}^{\text{Firm2}} \\ 0?$$

$$\overbrace{\underbrace{000} \underbrace{001} \underbrace{010} \underbrace{011} \underbrace{100} \underbrace{101} \underbrace{110} \underbrace{111}}^{\text{Firm3}} \\ 0?$$

2. If all three firms are mean to each other, will any one keep it the same way in the next period? That is, to check:


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

3. If all the firms are nice to each other in this period, will any one turn mean to the others in the next period? That is, to check:

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

4. If the other firms are nice to the host firm, who is mean to them, will the host

firm continuously take advantage of the others in the next period? That is, to check:



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 positions* can be considered *aggressive*. Since each strategy may have different answers, for each question 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 the 0.05 significance level, all of them are greater than 0.5. Therefore, loosely speaking, strategies evolved from our two sets of oligopoly games are not *gentle*, but *a little aggressive*. This result is not surprising and is consistent with the prevalence of predatory behavior observed.

Of course, not all strategies can help firms to survive well. To have a vivid picture of this, we took an average of the market shares of 30 shops owned by each oligopolist. If the game converged to a fixed point in action space, this average was taken by using the market share data from the last period. If the game converged to a periodic cycle, then the average was 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 share 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.

At this point, one may be tempted to ask: *What is the most competitive strategy?* In a *coevolutionary* context, however, this question is not really well defined, since everything depends on everything else. As a result, even though the strategy is exactly the same, depending on the strategies it competes against, the associated pay-off can be quite different. For example, both the third oligopolist 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 characteristics of the two extremes of strategies, i.e., the ones leading to great success (a 100% market share), and the ones leading to 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. To generalize, we ran a *regression* of the average

Table 10. Simulation results: attractors, case A

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

The sign '↔' refers to the start of another cycle.

market share on the frequency of '1' appearing in each cases, and the results are as follows:

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

In both cases, the regression coefficient is negative. Therefore, the regression results do suggest that cooperation 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 behavior of S_t given a specific set (ψ_1, ψ_2, ψ_3) , i.e., the *attractors* of S_t . Tables 10 and 11 display these attractors. For cases that do not have a unique attractor, we list all of them, with the one which was actually visited put in the first place. For example, case A1 has two attractors '000' and '100', but only '000' was realized as the final state of the game.

Two kinds of attractors appeared in the simulations. The first kind is *fixed points*. Interestingly, the *only* fixed point that was realized as the final state of the outcome is '000', i.e., *price wars* (cases A1, B1, B2, B6, B7, B8, and B9), which is quite understandable. If one firm decides to charge the low price from now on, then the only way for the other firms to survive is simply to 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., that firms would charge the high price at the same time. Case A6 is the only example. In this case, the cycle

Table 11. Simulation results: attractors, case B

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

The sign ' \leftrightarrow ' refers to the start of another cycle.

goes 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 one in two-dimensional state space: '00', '00', and '11'. In such a cycle, the two surviving firms would concurrently charge the low price for two periods, switch to the high price in the following period, and then switch back to the low price again. This behavior can be identified as the *nonaggression agreement* as discussed in Section 2. In case 6, this nonaggression agreement was a kind of emergent behavior; i.e., the agreement was not made at the beginning of the game. It was achieved in the complex path-dependent dynamics at the time when the first firm was already a big firm (with a 78% market share), and the third firm a medium-size one (with a 22% market share).

For the second kind of possibility, consider case A4. The attractor of case A4 has a periodic cycle with four periods. This cycle goes 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 charge both the high price and the low price twice, while at different times.

As we briefly mentioned above, 'charging high prices *in turn*' can be considered a collectively intelligent behavior in practice, because it is a *smart* way to avoid the violation of the anti-trust law. In this coordination, firms fake competition, while under the table they reach a tacit agreement to charge high prices half the time. By the tacit agreement, they are guaranteed not to lose market to each other.¹¹ Hence, this kind of behavior can also be classified as a type of *nonaggression agreement*. Moreover, since they are charging high prices at different times, it is very difficult to verify their anti-competition behavior. Therefore, 'charging high prices *in turn*' is an even more sophisticated design of cooperation. In addition to case A4, this type of emergent behavior is also observed in cases A3, A5, and A7. Therefore, *cooperative behavior does emerge in our simulations, but is manifested in a much more subtle way than in the IPD game.*¹²

¹¹ As indicated earlier, this 'tacit agreement' does not mean that agreement was made between firms, but that an apparent agreement evolved out of market interactions.

¹² Apart from not violating the anti-trust law, firms can benefit from this tacit agreement if both the sixth and the seventh conditions of the IPD game are violated, as discussed in Section 3.

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

$$\underbrace{000}_{\text{Price war}} \rightarrow 010 \rightarrow \underbrace{111}_{\text{Collusive pricing}} \leftarrow \quad (19)$$

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

$$001 \rightarrow 011 \rightarrow 100 \rightarrow 110 \leftarrow \quad (20)$$

is pretty much about the second stylized fact, i.e., that a dispersion of prices can persistently exist. As to the third stylized fact, i.e., that 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' behavior.

5. Concluding Remarks

In this study, GAs were applied to an oligopoly game. While firms (players) encounter a similar subtle decision on whether they should *defect* or *cooperate*, the Markov process characterization of firms' market shares makes the oligopoly game a nontrivial generalization of the IPD game. First, depending on the initial market shares and transition rules, the pay-off matrix is time-variant and state-dependent. Second, the inequalities which define the IPD game may be violated as time goes on, which means that even though an oligopoly game can satisfy the conditions of the IPD game in the beginning, the learning process randomly initiated may fail it in a later stage. Due to these extensions, one may 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 behavior) is not the dominant result as one may expect from the standard three-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 oligopolists' pricing behavior are captured by many of our simulations (cases A3, A4, A5, and A7). Furthermore, many of our simulations experienced a transition from an oligopoly industry to a monopoly (A2, A8, A9, A10, B3, B5) or duopoly industry (A6, B1, B4, B10). One thing that failed to be generated in our simulations is persistent collusive pricing. Persistent pricing wars instead seem to be the predominant outcomes. While a more advanced model could be attempted with GAs, we very much doubt if the rich nature of the oligopoly industry would change because of such sophistication.

Acknowledgements. 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.

References

- Abreu D, Pearce D, Stacchetti E (1986) Optimal cartel equilibria with imperfect monitoring. *Journal of Economic Theory* 39: 251–269
- Beaufils B, Delahaye JP, Mathieu P (1998) Complete classes of strategies for the classical iterated prisoner's dilemma. In Porto VW, Saravanan N, Waggen D, Eiben AE (eds.). *Evolutionary programming VII*. Springer, Berlin, pp 32–41
- Conte R, Hegselmann R, Terna P (eds) (1997) *Simulating social phenomena*. Lecture notes in economics and mathematical systems 456. Springer, Berlin
- Fudenberg D, Tirole J (1989) Noncooperative game theory for industrial organization: an introduction and overview. In Schmalensee R, Willig RD (eds). *Handbook on industrial organization*. North-Holland, Amsterdam, pp 259–328
- Green E, Porter RH (1984) Non-cooperative collusion under imperfect price information. *Econometrica* 52: 87–100
- Marks RE (2000) Playing games with genetic algorithms. In Chen S-H (ed). *Evolutionary computation in economics and finance*. Physica-Verlag, Berlin
- Midgley DF, Marks RE, Cooper LG (1997) Breeding competitive strategies. *Management Sciences* 43(3): 257–275
- Scherer FM, Ross D (1990) *Industrial market structure and economic performance* (3rd edn). Houghton-Mifflin, Boston, MA
- Shapiro C (1989) Theories of oligopoly behavior. In Schmalensee R, Willig RD (eds). *Handbook on industrial organization*. North-Holland, Amsterdam, pp 329–414
- Yao X, Darwen PJ (1994) An experimental study of n -person iterated prisoner's dilemma games. *Informatica* 18: 435–450

Author Biographies



Shu-Heng Chen received his Ph.D. from the Department of Economics at the University of California, Los Angeles, USA, in 1992. He is currently a professor of economics at National Chengchi University, Taipei, Taiwan, and the director of the AI-ECON Research Center. His research interests include the application of evolutionary computation to economics and finance.



Chih-chi Ni received her Ph.D. from the Department of Economics at National Chengchi University, Taipei, Taiwan, in 1999. She is currently an associate professor in the Department of Accounting at Chung Kou Institute of Technology and Commerce, Taipei, Taiwan. Her research interests include the application of evolutionary computation to economics.