

COMPETITIVENESS, COOPERATION, AND STRATEGIC INTERACTION

A classroom experiment on oligopoly

COMPETITIVIDAD, COOPERACIÓN E INTERACCIÓN ESTRATÉGICA

Un experimento oligopolístico en el aula

JOSÉ ANTONIO GARCÍA-MARTÍNEZ Jose.garciam@umh.es

CARLOS GUTIÉRREZ-HITA cgutierrez@umh.es

JOAQUÍN SÁNCHEZ-SORIANO joaquin@umh.es

Miguel Hernández University. Alicante. Spain

ABSTRACT

We run a classroom experiment on oligopoly with students enrolled on basic and medium level microeconomics courses. Students compete in a symmetric quantity setting environment. The experiment runs over an entire academic semester and is divided into 20 one-week rounds. We want to explore whether the effect of knowledge and social interaction between players modifies the cooperative and competitive behavior observed in similar experiments run in a lab. Our hypothesis is that players are socially influenced. Hence, individuals adjust behavior in a dynamic way aimed at maximizing profits, but also according to social pressures. Overall, we obtain different learning processes across academic levels and also slightly different behavior from that predicted by economic theory. We argue that students' utility function depends not only on profit levels but also on social relationships. Moreover, we believe that the effect of reputation plays an important role in our framework.

KEYWORDS

Cooperation; Economic experiments; Learning processes; Oligopoly Competition.

RESUMEN

Hemos llevado a cabo un experimento oligopolístico en el aula con estudiantes pertenecientes a los niveles básico y medio de la asignatura de microeconomía. Los estudiantes compitieron en un mercado donde todos tenían la misma función de costes ofreciendo cantidades de un bien en cada ronda. El experimento tuvo lugar durante un semestre académico completo y fue dividido en 20 rondas de una semana. Se quiere investigar si el efecto del aprendizaje y de la interacción social entre los jugadores modifica el comportamiento competitivo y cooperativo observado en experimentos similares llevados a cabo en un laboratorio. Nuestra hipótesis es que los jugadores están socialmente influenciados. Así pues, los individuos adecúan su comportamiento dinámicamente con el objetivo de maximizar beneficios pero también en concordancia con las presiones sociales. En general, se han observado diferentes procesos de aprendizaje por niveles académicos así como pequeñas diferencias de comportamiento con respecto a las pronosticadas por la teoría económica. Argumentamos que la función de utilidad de los estudiantes depende no sólo de los niveles de beneficio alcanzados sino también de las relaciones sociales. Más aún, pensamos que la reputación desempeña un papel relevante en el experimento de mercado que se ha llevado a cabo.

PALABRAS CLAVE

Competencia Oligopolística; Cooperación; Experimentos Económicos; Proceso de Aprendizaje.

INTRODUCTION

This paper addresses an interesting issue for experimental literature: Does economic background modify the behavior of players in an experiment? And also, do players behave differently when social interaction is allowed? If this is the case, players' strategies should be different from those that arise from a usual maximizing-profits game. Moreover, the learning processes towards standardized economic equilibria could be different. To highlight these issues we perform an economic experiment that takes place outside the computer lab. Most oligopoly experiments are run on computers (facilitating data collection and rapid implementation), but an experiment run manually or, in other words, by allowing real interaction between players in (and outside) class is interesting for a number of reasons. First, it is appealing for instructional purposes. Second, it gives us the opportunity to test how social concerns may influence economic strategies. Moreover, this approach has the added values of flexibility and not being dependent on the availability of computer rooms.

We run an oligopoly experiment with basic and medium level microeconomics students. A number of experimental subjects (players) play a game in which they compete in a symmetric quantity setting environment. The market is open for a whole semester divided into 20 rounds (one per week). In each background level players are divided into 7 markets of 10 players each. Thus, 140 students play the classroom experiment. Players decide their output at any time from Monday to Friday: There is no requirement to choose a strategy at any particular time over these five days. Under this environment, players can cheat each other before the strategy is chosen. Hence, it is possible for players to take into account not only economic information but also sociological concerns that may modify the final output decision. Moreover, we are interested in checking whether the concepts of adaptive learning and imitation are crucial in explaining the path towards equilibria in each economic background and across rounds. Our hypothesis is that sociological concerns are important and that, accordingly, players maximize a more general utility function which depends not only on profit-maximizing behavior but also on social pressures imposed by the group. We want to explore whether the effect of knowledge and the incentives to behave either competitively or cooperatively provide different market outcomes across different background levels and outcomes different from those predicted by economic theory. If players are influenced by social pressures,

*Acknowledgements: We wish to thank all participants in the 2nd Southern Europe Experimentalists Team (SEET) meeting at Agadir (Morocco) and especially Nikolaos Georgantzis for helpful and interesting comments. Carlos Gutiérrez-Hita and Joaquín Sánchez-Soriano acknowledge financial support from the Spanish Ministry of Education and Science (through research project MTM 2008-06778-C02-01) and Generalitat Valenciana (through research project ACOMP 2011-129). Carlos Gutiérrez-Hita also acknowledges financial support from the Basque Government (through research project IT-313-07). The usual disclaimers apply.

they adjust their behavior in a dynamic way accordingly. A consistent observation across markets is that neither best response dynamics nor partial adjustment to best response fully explains output fluctuations.

Literature on oligopoly experiments is abundant, and provides different conclusions depending on the assumptions made. An important issue is whether information about successful strategies of rivals may influence the learning process or, on the contrary, players follow an adaptive learning process to converge at any long run outcome. On one hand, Vega-Redondo (1997) points out that an evolutionary process in which firms can imitate their rivals' strategies yields convergence to the competitive equilibrium in a symmetric environment. On the other hand, Milgrom and Roberts (1991) propose a general setting known as adaptive learning to study dynamic interaction of subjects. Many papers have followed the latter approach to test the theoretical predictions of oligopoly models. Rassenti et al (2000) analyze some versions of Milgrom and Roberts' setting. In particular they study (i) best response dynamics; and (ii) partial adjustment to best response. In that paper experimental subjects play an asymmetric game. They seek to explore whether a repeated play by privately informed subjects converges to a unique, static, noncooperative Nash-Cournot equilibrium. It is found that this repeated interaction partially converges. Moreover, total output averaged over time periods and across experiments is greater than the predicted equilibrium total output. In our setting, we find that in both economic backgrounds players sooner or later converge to the predicted Walrasian output. A similar approach to Rassenti et al is found in Van Huyck et al (1990; 1994) but in these papers they attempt to select an equilibrium when there are multiple equilibria under homogeneous subjects. Huck et al (1999; 2000) test various learning theories in the context of a Cournot oligopoly by varying the information given to subjects. A crucial assumption here is that the learning and decision adjustment processes depend only on aggregate output of rivals in previous periods. The results show that some subjects imitate successful behavior if they have the necessary information, and if they do so then markets are more competitive. Other subjects follow a best reply process.

More recently, in Apesteguia et al (2007) a generalized theoretical approach is introduced to study imitation by testing different theories in a laboratory experiment by systematically varying information conditions. In their theoretical analysis they find that the predictions of previous imitation models are mainly explained by different information assumptions and to a lesser extent by different behavioral rules. Overall, the data provide evidence of imitation at individual level; moreover, individuals' propensity to imitate more successful actions is increasing in payoff differences. Apesteguia (2010) also tests the well-known result by Vega-Redondo (1997) that implies that imitation leads to the Walrasian outcome in symmetric Cournot oligopolies. This is exactly what our results show under Treatment I, in which there is no incentive to behave cooperatively. Apesteguia (2010) also shows that this result is not robust to the slightest asymmetry in costs and suggests a new theory based on a mix of imitation and other learning processes to explain the behavior of subjects better.

Social conflicts are also relevant in economic experiments. In our context, Treatment II allows for cooperative behavior by groups but also for individual profit-maximizing behavior. Thus, the subject has to balance these two opposite pressures in order to maximize his/her utility function. Our Treatment II is also related to group selection models. These models were developed in biology (Wilson, 1983), and some economic applications are outlined by Vega-Redondo (1993) and Sjöström et al. (1996). In group selection models, populations are fragmented into locally isolated groups with independent evolutions, corresponding to our markets. Within each group, there is some kind of individual selection or competition, but there is also selection or competition between groups. Thus, more efficient groups have an advantage over less efficient groups. The basic result is that group selection may outweigh the effect of individual selection. A similar result is found in our environment: it is better to promote internal cooperation than maximize individual profits when groups compete with one another. Along these lines, Blackwell (2003) reports on the results of a series of experiments designed to investigate the role of preferences on individual willingness to contribute to the provision of a group (excludable) versus a global (non-excludable) public good. Results show that when the average per capita return to society of the global public good exceeds the average per capita return to society of the group public good, individuals contribute more to the global good but do not reduce their contributions to the group public good. Choi (2007) investigates how altruism -benefiting fellow group members at a cost to oneself- and parochialism -hostility toward individuals not of one's own ethnic, racial, or other group- are common human behaviors. The intersection of the two, in his words parochial altruism, is puzzling from an evolutionary perspective because altruistic or parochial behavior reduces one's payoffs by comparison to what one would gain by eschewing these behaviors. But parochial altruism could evolve if parochialism promoted intergroup hostilities and the combination of altruism and parochialism contributed to success in those conflicts.¹

The rest of the paper is organized as follows. Section 2 presents economic theory predictions that arise in our environment. Section 3 defines the experimental environment. Section 4 shows and analyzes the results of the experiment. Section 5 concludes.

WHAT ECONOMIC THEORY SAYS

We run an experimental oligopoly comprising 10 players in each market. Players compete on quantities. Hence, each subject decides an amount of output in each trading round. The number of trading rounds is 20 and this is commonly known. Linear market demand is assumed with an inverse function,

$$P(q_{it}, q_{-it}) = 1000 - (q_{it} + \mathbf{1}' q_{-it})$$

¹ His game-theoretic analysis and agent-based simulations show that under conditions likely to have been experienced by late Pleistocene and early Holocene humans, neither parochialism nor altruism would have been viable singly, but by promoting group conflict, they could have evolved jointly.

where P is the market price determined as a function of the total output,

$$Q = q_{it} + \mathbf{1}' q_{-it}$$

$\mathbf{1}'$ is a vector of dimension 1×9 , q_{it} is the output of player i at round t , and q_{-it} is a vector which includes all output levels but i , in round t . We study a homogeneous multiperiod Cournot market. There are ten symmetric firms in each market at each level—basic and intermediate— (7 markets per level). Quantities could be chosen from a finite grid between 0 and 100 with 0.01 as the smallest step. The cost function for each economic subject is,

$$C_i(q_{it}) = \frac{q_{it}^2}{2} + 100q_{it} + 50$$

We assume a linear-quadratic cost function because it represents a large penalization as long as output increases. Thus, production and output levels close to the maximum capacity for a given subject i not only decrease the market price but also raise costs. As a result, the damage to profits is significant and subjects must take these two effects into account. By contrast, when costs are linear there is no reason to produce a low quantity and strategic interaction does not enhance the learning effect.

Below, we report the Nash-Cournot equilibrium, the Walrasian behavior (perfect competition) and the most cooperative equilibrium. We also include the eventual output in the case that a subject deviates from the most cooperative equilibrium.

When firms behave as Cournot competitors they maximize,

$$\max_{q_{it}} P(q_{it}, q_{-it})q_{it} - C_i(q_{it})$$

The Walrasian equilibrium holds when all ten firms take the price as given, so that they decide the amount of output to produce as a result of their cost structure. Notice that despite of firms maximize profits there is no strategic interaction from a game-theory point of view. Thus,

$$\max_{q_{it}} P \cdot q_{it} - C_i(q_{it})$$

The most collusive outcome (the monopoly outcome) is found when firms maximize joint profits (alternatively when they minimize total industry costs),

$$\max_{q_{it}} P(q_{it}, q_{-it})(q_{it} + \mathbf{1}' q_{-it}) - C_i(q_{it}) - \mathbf{1}' C_{-i}(q_{-it})$$

where $C_{-i}(q_{-it})$ is a vector which comprises the cost functions except for firm i .

Finally, if a given firm deviates, by taking the static optimal best response,

$$q_{it} = R_i(q_{-it})$$

evaluated at the most collusive output fixed by the others q_{-it}^* an optimal deviation strategy is found. Table 1 reports theoretic equilibrium values arising from our particular environment. π_{it} stands for profits earned by player i at round t .

Table 1.
Theoretical equilibria

Walras	Nash-Cournot	Collusion	Deviation
81.81	75	42.85	100.00
818.1	750	428.5	485.65
181.9	250	571.5	514.35
3,297.09	8,387.5	19,235.7	36,378.0

The next section explores the strategic interaction that arises when the classroom experiment is run.

EXPERIMENTAL ENVIRONMENT

An oligopoly experiment is run with basic and medium level microeconomics students, referred to hereafter as Background I and Background II, respectively. A number of players compete in a symmetric quantity setting environment. They interact 20 times in the course of a whole semester (once a week), in two different treatments of 10 times each. In each background level players are divided into 7 markets of 10 players each. Each trading round comprises five days per week, from 00.00 on Monday to 23.59 on Friday. Thus, players decide their output at any time from Monday to Friday: There is no requirement to choose a strategy at any particular time over these five days. By opening the market for five days we get a long time horizon in which subjects can interact and behave according to profit-oriented decisions but also to reputation/social concerns. Moreover, as some learning processes may take quite some time to converge (if they do so at all) we need enough trading phases. Under this environment, players can cheat one another prior to choosing their strategies. Hence, it is possible for players to take into account not only economic information but also sociological concerns that may modify the final output decision.

We run two types of experiment, which we refer to as Treatment I and Treatment II. Each is run for 10 rounds and players are fully informed (it is a show experiment): they know at the end of each round the strategies played by the rest of subjects, market prices, total output of the markets and firms' profits. This information is accumulated over the trials. Under this environment we want to test whether economic background significantly influences the behavior of players and also see what effect social interaction has on the choices of rivals.

Treatment I

Under this environment players take into account the profits obtained by their competitors in their market. The profits of each player in each round are accumulated so that at the end of the first ten rounds a final ranking is formed for each market. Thus, seven rankings are obtained for each background. Under this environment, a given player only needs to take into account the profits obtained by the other members of his/her own market. The payoff of a given player depends solely on what happens within his/her market. Within each market and on each level players are paid as follows: The best result is normalized to one and the worst is normalized to zero. The rest of the players are ranked by a linear combination between zero and one.

Treatment II

This differs from Treatment I in only one way: After ten rounds just one ranking of accumulated profits is formed with all the players from each background, i.e., players in all the markets belonging to Background I compete with each other, and the same goes for Background II. Thus, two rankings are obtained: One for Background I and another for Background II. Under this environment players need to take into account the results (profits) obtained by the competitors in their own market but they also compete with the rest of the players in markets on their background level.

Payoffs

In each treatment subjects are paid up to 0.5 extra points in their final qualifications. In each ranking of profits the winner receives 0.5 extra points and the worst qualified player none at all. The remaining payments are assigned following a linear combination between the profit of the winner and the profit of the worst qualified player. Under Treatment I each market has a winner. Under Treatment II each group of seven markets competes on its own economic level, so that there is only one ranking per economic level with only one winner in each. As each treatment's contribution to the final qualification is 1/2, the maximum extra score is one.

Finally, despite the number of players in each market is relatively large compared with experiments run on a computer (thus, a player would view his/her choice as not relevant

within the overall market process) we think that this is not a challenge in our environment, for two reasons. First, the market is open for a long time so information and interaction between players reinforce the idea of competition. Second, the cost function is linear-concave, so larger quantities yield higher production costs. As a result, we argue that players may learn in a relatively short number of trials.

RESULTS AND ANALYSIS

Below, the results of the experiment are decomposed from an aggregate environment approximation to a detailed view of market members' behavior.

Aggregate behavior: static and evolutionary

Table 2 below presents the most general approximation to the results in order to highlight the main features that characterize both static and evolutionary group behavior in Backgrounds I and II.

Table 2.
Statistical Analysis by Treatment and Group

Background	Group	Treatment I			Treatment II		
		Average	SD	Are agents' SD equal?	Average	SD	Are agents' SD equal?
I	A	82.63 ^w	8.54	YES*	72.75 ⁿ	18.97	YES*
	B	92.90 ^w	7.39	YES**	57.83	19.69	YES***
	C	82.01 ^w	11.08	YES**	79.18 ^w	15.55	NO
	D	77.24 ⁿ	15.30	YES***	59.93	17.45	YES***
	E	73.99 ⁿ	17.37	YES**	79.62 ^w	14.38	NO
	F	78.04 ^{n,w}	13.94	NO	55.98	15.27	YES***
	G	73.60 ⁿ	19.41	YES**	80.65 ^w	11.74	NO
II	A	78.70	9.36	YES**	53.73	16.29	YES***
	B	81.66 ^w	10.78	YES***	73.68 ⁿ	18.85	NO
	C	80.66 ^w	9.40	YES***	57.86	21.47	YES*
	D	81.17 ^w	7.08	YES***	50.34	15.60	YES***
	E	81.60 ^w	6.59	YES**	47.27	15.01	YES***
	F	80.78 ^w	5.87	YES**	61.23	20.86	YES***
	G	80.38 ^w	8.57	NO	53.43	17.53	YES**

w, *n* and *m*: There are no statistically significant differences between the mean of the group and the Walrasian solution, the Nash-Cournot solution and the monopoly solution, respectively (at 5% significance level). **p*-value ≥ 0.05 , ***p*-value ≥ 0.1 , ****p*-value ≥ 0.5 .

STATIC BEHAVIOR

We start by analyzing the results from an aggregate and static point of view in order to obtain a general view of the behavior of players under the experimental game detailed in Section 3. We first observe that there are slight differences between the average outputs of Treatment I and the average outputs of Treatment II, and between the average outputs of players with a higher-level background in economics and those with a lower-level background (see Table 2). In particular, in Treatment I the average outputs of the different groups are close to the Walrasian solution (9 out of 14) and to the Nash-Cournot solution (5 out of 14), while in Treatment II the predominant average behavior among all groups is not far from the monopolistic solution (9 out of 14). Several groups are also observed with average outputs close to the Walrasian solution (3 out of 14) and to the Nash-Cournot solution (2 out of 14). Therefore, in general, players behave more competitively in Treatment I and more cooperatively in Treatment II. One possible explanation for this could be that in Treatment I players compete with none another within the group without being worried about what is happening in the other groups, while in Treatment II players compete within the group but the results obtained by players in other groups affect them and they cannot control them. Therefore, there are two level of competition: internal and external. Perhaps the existence of rivals outside may lead to more cooperative behavior among the players in each group.

On the other hand, slight differences are also observed between the average outputs of the groups when we take into account the economic backgrounds of the players. In Treatment I the average outputs among those groups with a higher-level background in economics are close to the Walrasian solution (6 out of 7), while among those groups with a lower-level background they are equally divided between the Walrasian and Nash-Cournot solutions. In Treatment II the differences in average output between the groups with higher and lower-level economic backgrounds are bigger. The average output of the groups with higher-level backgrounds is not far from the monopolistic solution (6 out of 7), but that of the groups with lower-level backgrounds is more variable (3 out of 7 close to the Walrasian solution, 3 out of 7 close to the monopolistic solution and 1 out of 7 close to the Nash-Cournot solution). Nevertheless, we observe that in almost all cases the average outputs in Treatment I are more competitive than in Treatment II.

Likewise, it is also observed that the variability in the outputs is slightly less in Treatment I than in Treatment II. This is because players behave more homogeneously in Treatment I than in Treatment II. This may be because in Treatment I players collaborate less with one another, which leads them to vary their offers within a short range of potential outputs. However in Treatment II although players collaborate, from time to time they drastically deviate in order to obtain competitive gains, thinking about the final result; and this behavior leads to a higher variability.

Tabla 3
Linear trend of the standard deviation along the periods

Background	Group	Treatment I	Treatment II
		Correlation coefficient	Correlation coefficient
I	A	-0.51	0.90***
	B	0.04	0.10
	C	0.64*	0.56
	D	-0.95***	0.02
	E	-0.69*	-0.39
	F	-0.71*	-0.69*
	G	-0.07	0.73*
II	A	-0.78**	-0.57
	B	-0.76*	0.67*
	C	-0.72*	0.44
	D	-0.64*	-0.59
	E	-0.70*	0.80**
	F	-0.23	0.73*
	G	-0.70*	0.46

*p-value ≤ 0.05 , **p-value ≤ 0.01 , ***p-value ≤ 0.001 .

The last point of interest to be noted in the aggregate and static analysis is that several interesting results are observed in the analysis of the variability of the output of the players. The first is that in most groups there are no statistically significant differences either in Treatment I or in Treatment II between the standard deviations of the variable outputs of the players in the same group (see Table 3). This means that in general all players belonging to the same group use a similar range of units of production.

EVOLUTIONARY ANALYSIS

First, the evolution of the variability of the units produced in each round is analyzed, measured by their standard deviation. This variability is observed to decrease in the course

of the rounds in Treatment I (negative linear correlation) and increase in the course of the rounds in Treatment II (positive linear correlation), as shown in Table 3. This implies that players in Treatment I learn from round to round to respond to the strategies of their opponents, leading offers by all players to become closer and closer to one another and converge to the Walrasian solution. By contrast, in Treatment II it is observed that variability increases from round to round: it is argued that players either learn or follow strategic behavior. To some extent, players behave more cooperatively in early rounds (outputs not far from the monopolistic solution) whereas in the later rounds they deviate aggressively in order to protect themselves from the other groups and in order to obtain higher revenues in the final result.

The following two Figures (Figures 1a and 1b) show the aggregate behavior per economic level. Thus, for each round there are 70 observations.

Similar qualitative behavior is observed at both academic levels. In Treatment I the main difference is that at the second academic level players behave more aggressively from an economic point of view and also learn more rapidly. This means that players can be seen to approach standard economic theoretic equilibrium faster than at the first academic level. This occurs in each group in general. Of course, there are multiple profiles depending on the internal relationships and the trend in competition from round to round

Figure 1a.
Aggregate average behavior. Background I





Figure 1b.
Aggregate average behavior. Background II



within each group. After a sufficient number of rounds all players tend to behave as one might expect as a result of repeated economic interaction. We now study these general insights in greater detail.

The pictures above show particular behavior patterns. Picture 1 reports the aggregate mean, standard deviation and well-known theoretic equilibria: Walras, Nash-Cournot, and perfect collusion (the monopoly outcome). Picture 1(a) plots the aggregate output per round of each of the 70 players enrolled in Background I, regardless of what group they belong to. Picture 1(b) reports the same statistics and equilibria for Background II. In general terms, it is observed that outcomes of players enrolled in Background II on average go towards the Walrasian equilibrium faster than those in Background I. This occurs in Treatment I (rounds 1 to 10). In Treatment II (rounds 11 to 20) the same profile is observed, but now approaching some level of collusion (eventually, perfect collusion). In Background I, the average output in the first round is smaller than in the second one, and it takes five rounds to converge to the Walras output. However, Background II converges to it in just two rounds. On the other hand, in Treatment II the average output of Background I remains around Walras output for two rounds, then starts to decrease and finally increases a little. The Background II shows a similar pattern, but it decreases sooner and more.

Moreover, standard deviation in Treatment I is lower in Background II. These facts reveal that economic know-how may play an active role in determining not only the path to either Walras equilibrium or collusion, but also the homogeneity of players when they decide their outcomes. Picture 1 also reveals faster learning in Background II. In Treatment II perfect collusion is almost achieved in Background II. It is interesting to note that in some groups no-one deviates from collusion for a small number of rounds, so the monopoly outcome is achieved. By contrast, in Background I collusion is only partial. Cheating is observed at both economic levels but cheating from collusive paths is observed first and to a greater extent in Background II. This may also be caused by economic know-how. Finally, at both levels rational behavior is seen to emerge when players play a finite-repeated game: incentives to cheat increase as the end of the game approaches.

INSIGHTS PER ECONOMIC LEVEL: A COMPARISON

This subsection examines the behavior of each group within the two levels of economic background. Figure 2 plots the path followed by groups over the whole experiment. Figure 2(a) focuses on Background I and Figure 2(b) plots Background II. Notice that the first ten rounds correspond to Treatment I and the last ten rounds to Treatment II. They are described below to highlight the behavior described.

Figure 2a.
Average behaviour by markets, Background I





Figure 2b.
Average behaviour by markets, Background II



We focus first on Background I. First, the behavior of the different groups throughout Treatment I is compared. It is useful to calculate the average output in each round for each group.

In Treatment I there are two kinds of initial average behavior. Groups A, B and C start the experiment with high outputs, in particular two of them around the Walrasian output. By contrast groups D, E, F and G start the experiment with low average outputs, some of them close to the collusion output. They have a wrong perception of the game. In Treatment I, collusion does not offer any advantages. Payoffs depend exclusively on the position in the group profit ranking. However, they learn very fast. After four rounds, all groups are on average around the Walrasian output, which is the more competitive situation. The convergence to the Walrasian output is fairly clear for all groups, and it is noticeable how close the average behaviors of the groups are to that output. Some groups obtain negative profits for all members in one or two rounds when their outputs are too high. For example, group A in R3 and R5, group B in R1, group C in R6, group F in R6 and R8 and group G in R6. In all cases the next round has lower output levels with positive profits.

Standard deviation of output in each round and for each group is decreasing in general in Treatment I and the differences between them also decrease, with the exception of groups G and E, which have one or two players who choose outputs such as 50 even close to the last round.

In Treatment II, high levels of output can be observed in all groups in the first round (R11). They have not yet learnt that they need to cooperate and that collusion is their best chance to get a high position in the global profit ranking. They need some rounds to fully understand that a new treatment is working. In the second round (R12), all groups except group D are still around the Walrasian output. Group D obtains an average output between Nash-Cournot and collusion output. Observation of the player outputs in group D reveals that six players choose outputs between 55 and 49. Although they have not been taught about oligopoly and collusion and are unable to calculate the collusion output, it is clear that they try to collude. However, there are members of group D who still have high outputs and consequently get high profits. The following round is more competitive for group D and the average output increases to the Nash-Cournot output. After that, the average output is decreasing round after round until it almost reaches the collusion value. In this group the cooperative players do not change their behavior if any member deviates and chooses a higher output. In R18 all members of group D except one choose an output of 43.

In Treatment II the learning process is more complex. Some degree of cooperation is needed. Thus, a more heterogeneous behavior from group to group is found. There are groups that never try to cooperate, such as groups G and E; surprisingly, these groups started Treatment I almost in collusion. In group C, too few players try to collude for just a few rounds, and the majority of the group always choose very high outputs. In this case, the average goes down to the Nash-Cournot value and then goes up. It ends up a little higher than the Walrasian output. Group F colludes after two rounds and B after three and then both groups maintain collusion. The winner of Treatment II is a member of group B, and he cooperates but deviates three times. Group B is very cooperative and not too sensitive to deviations in the sense that cooperators do not change their behavior if a player deviates, choosing a high output and getting much higher profits than them.

All members of group A except one collude to 48 after four rounds (R15). There is a player who always chooses 81. In R16, there are two players who do not collude. In R17, these two are joined by another player who also does not collude. In this round, the cooperating members change their output from 48 to 43. They have learnt a better output level for collusion. In round 18 there are only three agents producing 43. Eventually, just one player chooses 43 the rest have extremely high outputs. In this group, the members are sensitive to deviating players. This group would perhaps have been able to collude for more rounds if there had not been this one particular player who always chose 81.

The ranking of the groups by the aggregate profits obtained in Treatment II is: F, D, B, A, C, E and G. The standard deviation in this second treatment is more heterogeneous. There are some groups, such as F and B, in which collusion is so perfect that the standard deviation is zero or very close to zero. In general the standard deviation shows greater differences between groups than in Treatment I. The rise and fall of the values show the tension between cooperation and deviation. When some players deviate from collusion and choose high output levels the standard deviation increases.

It is important to point out again that players are anonymous. However sooner or later they find out who is who because the experiment lasts twenty weeks. If they know each other, the commitment to cooperation is greater. They have to see each other every day in the classroom. There is social pressure. At that point, the behavior of a group also depends on the relationship between its members. That could explain the different behaviors among groups in Treatment II with the same economic background.

In conclusion, the learning process is faster and more homogeneous among groups in Treatment I than in Treatment II because there is no need of cooperation. It is noticeable how fast the groups converge to Walrasian output and how small the difference between them is. However, in Treatment II more differences between groups are found. Some groups are relatively successful in achieving cooperation, others try but fail, and others do not even try to cooperate.

We now look at Background II group by group. In Treatment I the learning process is fairly fast compared with Background I. In Treatment II some markets get almost perfect collusive profits. This is the main difference with Background I. However, once some degree of collusive profits is achieved, some players deviate from the collusive path in early rounds as compared with Background I. However they frequently come back to (almost) collusive outcomes. These two observations can be explained by differences in economic background but we argue that there is also a sociological effect. As individuals have known one another for two academic years instead of one (as in Background I), social pressures may be reinforcing cooperative behavior. Overall, 3 types of group can be distinguished: First, a group comprising markets A, D, E, and G. Their common feature is that they converge under Treatment I to Walras in a small number of rounds and almost perfectly. Under Treatment II, they achieve perfect collusion in three rounds. Moreover, two profiles can be distinguished within this group. First, groups A and D maintain perfect collusion with only a few deviations and finish the game in collusion. By contrast, groups E and G break perfect collusion when they approach the end of the game: they finish with average outputs above perfect collusion with some players opting for high quantities (approaching 100 units of output). These differences can be explained by differences in social cohesion within each group. As social relationships between players increase the incentives to cheat can be expected to decrease because of subsequent social damage. Surprisingly, this occurs although the game is finite.

The second group comprises markets B and F. In Treatment I they finish at Walras but they approach it with some noise (the standard deviation significantly increases in some rounds). In Treatment II they approach Nash-Cournot equilibrium with some players offering monopoly outcomes while others behave competitively by choosing high levels of output. Summing up, dispersion is greater than in the first group. This reveals a low level of coordination and, perhaps, the existence of various subgroups or a lack of leadership. The third group comprises market C alone. This market approaches Walras with noise in Treatment I and also tries to collude in Treatment II, achieving some degree of collusion but with a high dispersion of players' outcomes. Under both treatments in this market agents behave heterogeneously, drawing a random output path over the 20 rounds.

To highlight the most important observations, consider the market that fits the theory best and that which fits it worst. Market D behaves in accordance with economic theory in both treatments. Players can be seen to converge to a Walras equilibrium in Treatment I with a lower dispersion as the number of rounds increase (as standard deviation indicates). In particular, the mean is between 70 and 90 units of output around Treatment I. By round 3 of Treatment II they converge to perfect collusive behavior with only one deviation in round 5. This group finishes Treatment II in perfect collusion, which reveals a great social cohesion or sufficient social pressure to prevent players from cheating on the cartel agreement. It is noticeable that such behavior occurs only in Background II. By contrast market F is more eclectic. Under Treatment I its behavior is quite similar to the rest of the groups, except for some erratic outcomes in the middle rounds; in Treatment II it proves to have extremely low levels of coordination. From rounds 3 to 5 players approach almost perfect collusion but from round 6 onwards many players cheat by choosing higher outcomes. It seems that economic know-how prevents this group from obtaining negative profits; thus, it approaches the Nash-Cournot outcome instead of the Walras outcome.

CONCLUSIONS

The main findings of this experimental study can be summed up as follows. Taken alone, the most aggregate average outputs of each market (see Table 2) in general are not far from the Nash-Cournot, Walrasian and monopolistic solutions. This is observed especially for the cases of the Walrasian and Nash solutions (see Table 2). This is because the monopolistic solution is more sensitive to the unilateral deviation of players, for two basic reasons: First, it is difficult to get the whole group to agree to behave collusively. Second, even if such an agreement is achieved it is difficult to maintain over time, taking into account that players also face a situation with a finite horizon, which favors the deviation of players in the last rounds of the game.

Competitive behavior is detected in Treatment I and collusive behavior in Treatment II. This is because in Treatment I there is only internal competition, while in Treatment II there is both internal and external competition. The external competition is an externality which provokes a collaborative reaction in the members of a group in order to defend themselves against external competitors.

There are differences between average outputs depending on the background in economics. The students with Background II have already undergone part of the learning process before the experiment starts. We observe that players learn throughout the ten rounds in Treatment I. The variability in players' outputs decreases over the ten rounds, which implies that players converge more or less to the same output. This in turn means that there is an action-reaction effect which favors the learning of players, leading (in the long term) towards an equilibrium solution. We also observe that players behave strategically in Treatment II, because the variability of the outputs increases over the ten rounds.

This, in some sense, implies that the players collude at the beginning and deviate in the last rounds to obtain a competitive advantage for the final result.

When the data are observed in greater detail (see Pictures 2(a) and 2(b)) the evolution of the average outcome of each market in each round emerges. In Treatment I markets converge to the Walrasian solution faster in Background II than in Background I. In Background I, some groups start Treatment I with collusive behavior but in a few rounds they are close to Walrasian output. It is noticeable how fast the groups converge to Walrasian output and how little difference there is between them. In Treatment II the learning process is more complex. They need some degree of cooperation. Thus, more heterogeneous behavior is found among groups. Some groups are relatively successful in achieving cooperation, others try to cooperate but fail, and others do not even try. The students from Background II are in general more successful in cooperating than the students from Background I. This may imply that their level of knowledge influences the behavior of players. We also argue that the effect of social pressures is more evident in Background II: cooperative behavior is more likely to occur, but players also come back to the *group discipline* in only one round in the event of deviations. The present experiment provides new insights and suggests that results from lab experiments should be checked in a more *open environment* where players can interact with one another.

REFERENCES

- Apesteguia, J., Huck, S., Oechssler, J. 2007. "Imitation: Theory and experimental evidence." *Journal of Economic Theory* 136:217-235.
- Apesteguia, J., Huck, S., Oechssler, J. 2010. "Imitation and the evolution of Walrasian behavior: Theoretically fragile but behaviorally robust." *Journal of Economic Theory* 145:1603-1617.
- Backwell, C., McKee, M. 2003. "Only for my own neighborhood? Preferences and voluntary provision of local and global public goods." *Journal of Economic Behavior and Organization* 52: 115-131.
- Choi, J.K., Bowles, S. 2007. "The Coevolution of Parochial Altruism and War." *Science* 318: 636-671.
- Huck, S., Normann, H., Oechssler, J. 1999. "Learning in Cournot oligopoly: An experiment." *Economic Journal* 454:80-95.
- Huck, S., Normann, H., Oechssler, J. 2000. "Does information about competitors' actions increase or decrease competition in experimental oligopoly markets." *International Journal of Industrial Organization* 18:39-58.
- Milgrom, P., Roberts, J. 1991. "Adaptive and sophisticated learning in normal form games." *Games and Economic Behavior* 3(1):82-100.
- Rassenti, S., Reynolds, S., Smith, V.L. 2000. "Adaptation and convergence of behavior in repeated experimental Cournot games." *Journal of Economic Behavior and Organization* 41:117-146.

- Sjöström, T., Weitzman, M. 1996. "Competition and the evolution of efficiency." *Journal of Economic Behavior and Organization* 30:25-43.
- Van Huyck, J., Battalio, R., Beil, R. 1990. "Tacit coordination games, strategic uncertainty, and coordination failure." *American Economic Review* 80(1):234-248.
- Van Huyck, J., Cook, J., Battalio, R. 1994. "Selection dynamics, asymptotic stability, and adaptive behavior." *Journal of Political Economy* 102:975-1005.
- Vega-Redondo, F. 1997. "The evolution of Walrasian behavior." *Econometrica* 65(2):375-384.
- Vega-Redondo, F. 1993. "Competition and culture in an evolutionary process of equilibrium selection: A simple example." *Games and Economic Behavior* 5:618-631.
- Wilson, D. 1983. "The group selection controversy." *Annual Review of Ecology and Systematics* 14: 159-187.

JOSÉ A. GARCÍA-MARTÍNEZ teaches at the Departamento de Estudios Económicos y Financieros [Dept. of Economic & Financial Studies] at Universidad Miguel Hernández. He holds a PhD in Quantitative Economics from Universidad de Alicante. His research interests include game theory, behavioral economics and experimental economics. He has published in *Journal of Economic Theory*.

CARLOS GUTIÉRREZ-HITA got his Bachelor degree in Economics at University Complutense of Madrid. After his Master in Quantitative Economics at University of Alicante he completed his Ph.D in Economics at University Jaume I. He is with the department of Economics and Financial Studies at Universitas Miguel Hernández where he teaches Microeconomic Theory and Industrial Economics. His research interest is focused on industrial organization and, in particular, electricity markets, applied game theory and regulatory policy.

JOAQUÍN SÁNCHEZ-SORIANO received a Mathematics Science Degree in 1991 and a PhD in Mathematics Science in 1998 both from Murcia University (Spain). He was Director of the Center of Operations Research of the University Miguel Hernandez of Elche (Spain). He is Full Professor in the Department of Statistics, Mathematics and Computer Science of the Miguel Hernández University. He is author or co-author of more than 60 research articles published in journals and books. His research interests include operations research and game theory and their applications, in particular to radio resource management, logistic problems and market design.

RECEIVED: 14 July 2011

ACCEPTED: 28 November 2011

APPENDIX

INSTRUCTIONS FOR STUDENTS²

This is an experiment in the economics of market decision-making. There are 7 independent markets known as A, B, C, D, E, F, and G.

In this experiment you will participate in the B-market, in which you will be producing and selling units of a fictitious good. In this B-market, there will be 10 producers and you will be known as producer B5. You are free to produce and sell as many units as you wish in the interval, $q \in (0,100)$, where q is the level of output that you produce. All the output produced will be sold. However, the more that is produced the less will be paid per unit of output.

This experiment is divided into two treatments of 10 rounds each. A round takes a week. You must select the quantity of the good that you will produce and send it by E-mail before the weekend. You are free to change your output/sales quantity once each week. For example, if you choose to produce 50 units, you must send B5-50 any day from Monday to Friday (inclusive). Each other producer from the B-market will also make a quantity decision. Q is the sum of the quantities chosen by all other producers in the B-market. The total quantity offered for sale (and bought by the buyers) each week is therefore $q + Q$.

The inverse demand function of your B-market is in any round:

$$P(q,Q) = 1000 - (q + Q)$$

Your cost function is in any round:

$$C(q) = \frac{q^2}{2} + 100q + 50$$

Therefore, your profit function is in any round.

$$\Pi(q,Q) = P(q,Q)q - C(q) = (1000 - (q + Q))q - \left(\frac{q^2}{2} + 100q + 50\right)$$

For example, if the outputs of a week are:

$$B1 = 50; B2 = 24; B3 = 56; B4 = 67; B5 = 34;$$

$$B6 = 68; B7 = 56; B8 = 90; B9 = 15; B10 = 21$$

$$Q = 50 + 24 + 56 + 67 + 34 + 68 + 56 + 90 + 15 + 21$$

$$\text{Price: } P(q,Q) = 1000 - (q + Q) = 1000 - (481) = 519$$

$$\text{Profits of producer B5: } \Pi(q,Q) = (519)34 - \left(\frac{34^2}{2} + 100(34) + 50\right) = 13618$$

² The instructions reproduced here are those given to the student labeled B5 in market 5. Each student received the same instructions: only the market and the player number are changed.

All producers in all markets have the same costs and face the same market demand function, so the profit function is the same across players and markets. At the end of the week (after Friday), you will receive an E-mail with the following information about the B-market in that week: The weekly outputs of each producer, the sum of all weekly outputs, the price of the good per unit in that week, the weekly profit made by each producer, and the accumulated profits that you and the others producers have earned up to this round. In each round you can choose either the same output as in the previous round or a different one.

After 10 weeks, a ranking of your market (B-market) will be drawn up, with the accumulated profits. Your reward for participating in this first stage will depend on that ranking. The seller from your market with the highest accumulated profit (called Π^{TOP1}) will get 0.5 extra points in the final mark. On the other hand, the seller from your market which the lowest accumulate profit (called Π_{LOW1}) will get zero extra point in the final mark. If you get an accumulate profit between them, your note will depend linearly on the distance to both extremes according to the following function where Π is your accumulate profits:

$$FinalMark1 = \frac{1}{2} \left(\frac{\Pi - \Pi_{LOW1}}{\Pi^{TOP1} - \Pi_{LOW1}} \right)$$

After 10 weeks a new stage will start. You will receive information about the functioning of the experimental game at the beginning of round 11.

SECOND TREATMENT

There is only one difference with Treatment I. After 10 rounds (for rounds 11 to 20) just one ranking will be drawn up with the accumulated profits of all the producers from the seven markets. Your reward for participating in this treatment will depend on that ranking. The seller with the highest accumulated profit (called Π^{TOP2}) will get 0.5 extra point in the final mark. On the other hand, the seller which the lowest accumulate profit (called Π_{LOW2}) will get zero extra points in his/her final mark. If your accumulated profit is between the two, your mark will depend linearly on your distance from both extremes according to the following function, where Π is your accumulated profit:

$$FinalMark2 = \frac{1}{2} \left(\frac{\Pi - \Pi_{LOW2}}{\Pi^{TOP2} - \Pi_{LOW2}} \right)$$