Price competition between teams[1]

Gary Bornstein* and Uri Gneezy**

August 30, 1998

Abstract: Economic agents (e.g., firms, corporations) are often treated as unitary players. The internal organization of these agents and, in particular, the possibility of conflicting interests within agents, are overlooked. The present study uses an experimental approach to examine whether the market is sensitive to the internal structure of the competitors. Toward this goal, we modeled the competitors in a price competition duopoly game as three-player teams. Each player simultaneously demanded a price and the team whose total demand was lower won the competition and was paid its price. The losing team was paid nothing. In case of a tie, each team was paid half its price. This composite duopoly was studied under two conditions; a cooperative treatment in which the team's profit was divided equally amongst its members and a non-cooperative one in which each individual member was paid her own price. Whereas the unique Nash equilibrium is for each player in either treatment to demand the minimal amount possible, we predicted, based on reinforcement learning, that convergence to the competitive price would be much faster in the cooperative treatment than in the non-cooperative one. The experimental results firmly confirmed this prediction.

* Department of Psychology and the Center for Rationality and Interactive Decision Theory, The Hebrew University of Jerusalem.
** Faculty of Industrial Engineering, Technion – Israel Institute of Technology, Haifa.

[1] This research was supported by a grant from the Israel Science Foundation (1997-2000) and is part of the EC-TMR ENDEAR project (1998-2001). Please address correspondence to Gary Bornstein, Department of Psychology, the Hebrew University, Jerusalem 91905, Israel. E-Mail: MSGARY@mscc.huji.ac.il

Price competition between teams

1. Introduction

One of the most prevalent assumptions in economic theory is that the agents (e.g., firms, corporations) operating in the market can be treated as unitary players. The internal organization of these agents and, most importantly, the possibility of conflicting interests within agents, are dismissed as "inessential complexities" (Davis & Holt, 1993). This simplifying assumption is reflected in experimental markets as well, where "firms" are commonly represented by individual subjects.

The purpose of the present experiment was to examine whether the market is sensitive to the internal structure of the competing agents and, in particular, to the existence of internal conflict. To motivate this line of inquiry consider, for example, a recent competition between Boeing and Airbus over a large deal with El Al Airlines. To simplify, assume that the airline's decision was based solely on price. This market can be conceptualized as a simple duopoly treating the two competitors as unitary players. Alternatively, one can consider a more complex (and probably more realistic) market structure where the two competitors consist of coalitions of firms, each developing and/or producing a different part of the airplane (e.g., engine, avionics, etc.). Each firm in the coalition sets its price independently and the price of the final product is the sum of prices demanded by the individual firms. Conflict or competition within each coalition can stem from the fact that, while all the members in a coalition have a common interest in setting a competitive price and winning the competition, each individual member also has an interest in maximizing its own share in the group's profit.[2]

The market in our experiment was operationalized as a competition between two teams with three players in each team. Each player simultaneously demanded a price (an integer between 2 and 25). The team whose total demand was smaller won the competition and was paid its price.[3] The losing team was paid nothing. In case of a tie, each team was paid half of its demanded price. This team game (Palfrey & Rosenthal, 1983; Rapoport & Bornstein, 1987; Bornstein, 1992) was played under two conditions that differed with regard to the way in which the team's profit was distributed among the three team members. In the non-cooperative treatment each individual player was paid her own demand if the team won, and half of her demand if the game was tied. In the cooperative treatment the team's profit for winning or tying the game was divided equally among its members. Thus, if the team won the competition, each member was paid one third of the total price asked by her team. And if the game was tied, each player received one sixth of the team's demand.

Predictions: The unique strict Nash equilibrium of the above team game, regardless of how profits are divided within the teams, is for each team, and hence for each team member, to demand the minimal price (i.e., 2 points per player). Thus, if one assumes that the players are perfectly rational, one should expect them to exhibit the same equilibrium behavior in both treatments.

However, if one considers the adaptive behavior of goal-oriented players who are not fully rational, and therefore, cannot be expected to play the equilibrium strategy right from the start, the difference between the two treatments becomes

[2] Alternatively, one can conceive of each coalition as consisting of different segments (e.g., management, workers, and stock-holders) within the same firm.

apparent. When behavior is away from equilibrium, that is, when prices are above the competitive minimum, the non-cooperative treatment provides each player with an opportunity, indeed a temptation, to free ride. Namely, if the other players in her team settle for a low price, player *i* can demand a higher price and still, possibly, win. In the cooperative treatment, on the other hand, where equal division of profits is imposed, the opportunity for free riding is eliminated.

To illustrate this point, assume that the members of team B make their decisions while knowing that the members of team A have already demanded a total of k points. The best response for team B in this case is to demand a total of k-1 points. In the non-cooperative treatment each individual player prefers that her team-mates demand as little as possible, while she demands the rest. At the extreme, player *i* prefers that each of her team-mates demand the minimum of 2 points, while she asks for the remaining k-5 points. The internal game in this treatment is thus a game of chicken (Bornstein, Budescu, & Zamir, 1997). In the cooperative treatment, on the other hand, team members are indifferent as to how the total of k-1 is reached, since any combination of individual prices, as long as they sum up to k-1, results in a payoff of (k-1)/3 for each member. The internal game in this treatment is thus a pure coordination problem.

This difference in the internal payoff structure is a key feature of our design. It enables us to vary the nature of the team while keeping the number of decision-makers in each team and the outcome of their joint decision vis-a-vis the other team intact. In the cooperative treatment there is no conflict of interests within teams and, therefore, each team can be considered a unitary player. In the non-cooperative

[3]For an experimental study of this game using unitary players, see Dufwenberg and Gneezy (2000).

treatment, where conflict exists within the teams as well, the unitary player assumption no longer holds.[4]

The game in the present experiment was played recurrently for 100 rounds. Generally speaking, we expect that the participants, as they gain more experience with the task, will be more likely to choose the equilibrium strategy of the stage-game. However, we predict that prices will decline much faster in the cooperative treatment than in the non-cooperative treatment. The rationale behind this prediction is rather intuitive. In both treatments a high demand by player *i* is likely to result in her team losing the competition and *i* receiving a payoff of zero. However, if *i*'s team ends up winning the competition, her payoff is much higher in the non-cooperative treatment, where she is paid her asked price in full, than in the cooperative treatment, where she must share it with her team-mates. This difference in the incentive structure is bound to result in different learning dynamics. Since high demands are punished more consistently in the cooperative treatment, reinforcement learning would impel subjects faster towards the stage-game equilibrium in this treatment than in the non-cooperative treatment.

2. Experimental Procedure

<u>Subjects and Design</u>: The participants were 120 undergraduate students (86 females and 34 males) at the Hebrew University of Jerusalem with no previous experience with the task. Subjects were recruited by campus advertisements offering monetary rewards for participating in a decision task. Subjects participated in the

[4] As argued by Luce and Raiffa (1957) any group of decision makers "which can be thought of as having a unitary interest motivating its decisions can be treated as an individual in the theory" (p. 13).

experiment in cohorts of 12; five cohorts took part in the non-cooperative treatment and five in the cooperative treatment.

Procedure: Upon arrival each participant received NIS 10 for showing up and was seated in separate cubicle facing a personal computer. The subjects were given written instructions concerning the rules and payoffs of the game (see Appendix) and were asked to follow these instructions while the experimenter read them aloud. Then subjects were given a quiz to test their understanding. Their answers were checked by the experimenters and, when necessary, explanations were repeated. Subjects were also told that to ensure the confidentiality of their decisions they would receive their payment in sealed envelopes and leave the laboratory one at a time with no opportunity to meet the other participants.

Subjects played 100 rounds of the game. The number of rounds to be played was made known in advance. At the beginning of each round the 12 subjects were randomly divided into three-person teams and each team was matched randomly with another team. This random-matching protocol was carefully explained to the participants. In each round each subject had to enter a demand of between 2 and 25 points. Following the completion of the round, the subject received feedback concerning (a) the total number of points demanded by the three members of her group in that round; (b) the total number of points demanded by the three members of the other group with which her group was matched; (c) the number of points she earned in this round; and (d) her cumulative earnings (in points).

Following the last round, the participants were debriefed on the rationale and purpose of the study. The points were cashed in at a rate of NIS 1 per 10 points (1 New Israel Shekel was equal to about $0.30 at the time the experiment took place) and the participants were dismissed individually.

3. Results

Overall means: First, we look at price demands averaged across the 100 periods and 5 cohorts in each treatment. The mean demand per period in the non-cooperative treatment was 8.68 points, as compared with 5.43 points in the cooperative treatment. To test this difference for statistical significance we calculated the mean demand per period for each 12-person cohort separately and analyzed the 10 means using a Wilcoxon non-parametric rank test. As can be seen in Table 1, the means are perfectly ordered in the sense that all five means in the non-cooperative treatment are larger than the five means in the cooperative treatment. The statistical test is, of course, significant ($z = 2.51$, $p < .012$). The same rank-ordering, and consequently the same statistical result, hold for the mean winning price and the mean profit, which also appear in Table 1.

Finally, we computed the correlation between individual demands and individual profit in both treatments. In the cooperative treatment this correlation was significantly negative (Pearson $r = -.258$, $p < .05$), whereas in the non-cooperative treatment the correlation was essentially zero ($r = .078$, *ns*). This finding directly supports our assertion that high individual demands are more reliably associated with lower payoffs in the cooperative treatment than in the non-cooperative treatment.

<Insert Table 1 about here>

Convergence: The mean price per round in each treatment appears in Figure 1. As can be seen in this figure, subjects' initial prices in the two treatments were almost identical. The mean number of points demanded in the very first round was 11.8 and 12.27 in the non-cooperative and cooperative treatments, respectively. These means are not statistically different ($t_{(118)}$=.441, ns). However, before long, subjects in the cooperative treatment began reducing their bids as compared with those in the non-cooperative treatment. The mean price in the first 10-round block was already almost 2 points lower in the cooperative treatment than in the non-cooperative treatment (M=8.55 and M=10.36, respectively). This difference is statistically significant by a Wilcoxon rank-test (z = 1.67, p < .05, one-tail), again using the 12-person cohort as the unit of analysis. The difference in first-block winning prices between the two treatments (M=9.08 and M=7.09 in the non-cooperative and cooperative treatments, respectively) is also significant (z=2.09, p<.05).

The difference between the two treatments became much more pronounced as the game progressed. In the last 10 rounds (10th block) the average price in the non-cooperative treatment was 7.92 points -- more than twice the 3.42 points demanded on average in the cooperative treatment. This difference is significant by a Wilcoxon test using the 12-person cohort as the unit of analysis (z=2.51, p<0.2). The same statistical result holds for the mean winning bids (M=6.93 and M=2.76 in the non-cooperative and cooperative treatments, respectively). In fact, at this final stage of the game the mean winning price in the cooperative treatment is quite close the competitive equilibrium (i.e., 2). The mean asked price and the mean winning price for each cohort appears in Table 2.

<Insert Table 2 about here>

Figure 2 plots the mean price per round separately for each 12-person cohort. As can be seen in this figure, the general pattern of results described above is characteristic of the individual cohorts as well. In the non-cooperative treatment, prices decline very slowly, and occasionally even rise, whereas in the cooperative treatment prices decline much faster to approximate the competitive price. The only exception seems to be cohort #4 in the non-cooperative treatment where demands decreased substantially over time. But even in this particular cohort, prices remained relatively high and never quite reached the competitive price. In fact, a look at the last block data in Table 2 will confirm that, while the mean demand in cohort #4 is indeed the lowest in the non-cooperative treatment, it is still higher than that in any of the cohorts in the cooperative treatment.

Learning: To examine the learning process more formally, we employed the reinforcement learning model as formulated by Roth and Erev (1995) and Erev and Roth (1998). The basic principle underlying the Roth and Erev model is the "Law of Effect" (Thorndike, 1898) which states that choices that have led to good outcomes in the past are more likely to be repeated in the future.

In applying the model to our price competition game, we assumed that Player i, when deciding how to act, considers the $K = 24$ pure strategies (i.e., choosing an integer between 2 and 25). At time t=1, before any experience with the task has been acquired, Player i's initial propensity to play each of the K strategies is evenly distributed. Her propensity to select a particular strategy, say strategy k, at time t=1 is denoted by $q_{ik}(1)$. If Player i plays strategy k at time t and receives a payoff of x, then the propensity to play k is updated by setting $q_{ik}(t+1) = q_{ik}(t) + x$, while for all other pure strategies j, $q_{ij}(t+1) = q_{ij}(t)$.

The probability $p_{ik}(t)$ that Player i will play strategy k at time t is $p_{ik}(t) = q_{ik}(t)/\Sigma q_{ij}(t)$, the ratio of the propensity to play the kth strategy divided by the sum of the propensities to play each of the 24 strategies. We assumed that at t = 1 all players have the same initial propensity. In setting this initial propensity we considered two factors: the ratio $q_{ik}(1)/q_{ij}(1)$ of the propensities to play strategy k and not k, which determines the probability that k will be played at time t = 1; and the sum of the initial propensities over the 24 strategies, $S(1) = q_{ik}(1)+q_{ij}(1)$. Following Roth and Erev, we set $S(1) = 10$.[5]

We derived the predictions of the Roth and Erev (1995) learning model by having virtual subjects play each other in a simulation of the experimental environment. Figure 3 presents the mean bid in 200 simulated groups, each playing 100 rounds of the game under either the non-cooperative or cooperative treatments. As can be seen in the figure, the predictions of the learning model are clearly in line with our experimental findings. In both treatments the simulated players, like the actual ones, learned to decrease their demands over time, however, they learned to do so much faster in the cooperative treatment than in the non-cooperative treatment.

<Insert Figure 3 about here>

5. Discussion

The field of industrial organization focuses on the relationship between trading institutions and market performance. Issues central to this research are competition, collusion, and efficiency in price-setting (Bertrand) and quantity-setting (Cournot) environments. Particular theoretical and empirical attention has been given to the effects of market concentration on competition. The number of competitors and their respective "market power" have a direct influence on the market's outcomes, and it is typically assumed (e.g., as reflected in antitrust policies) that more competition results in more "efficient" markets.

The present paper explored the relation between competition and performance

[5] This factor can be thought of as the *strength* of the initial propensities. When the value of $S(1)$ is large, the initial propensities are strong and learning is relatively slow. When the value of $S(1)$ is small, the initial propensities are weak and adaptation occurs more quickly.

from a different angle. Rather than increasing the number of agents competing in the market, we introduced competition *within* each agent, and studied the effect of this internal competition on the market's performance. The classic model of price competition (named after Bertrand, 1883) is mute to our manipulation. The model predicts that prices (even) in a duopolistic market will be equal to the marginal cost, and each duopolist, whether a unitary actor or a coalition of actors, will demand zero profit in equilibrium (e.g., Tirole, 1994). In the context of our experiment, this means that the unique Nash equilibrium in both treatments is, for each team, and hence for each individual player, to demand the minimal price (i.e., 2 per player).

Nevertheless, based on reinforcement learning principle we predicted that convergence to the competitive price will be much faster in the cooperative treatment where the teams are free of internal conflict (and, thus, can be regarded unitary players) than in the non-cooperative treatment where free riding within the teams is possible. Our experimental results strongly confirm this prediction. At least in the intermediate level, the effect of experience seems to depend on features of the environment different from those which determine equilibrium (Erev & Roth, 1998).

Since the primary interest of the present experiment was in studying behavior in the single-shot (i.e., stage) game, we imposed treatments of play that prevent the use of repeated-game strategies. Subjects participated in the experiment in cohorts of 12. Each cohort played a prespecified number of periods to ensure that the set of equilibria is not increased, so that the recurrent game has the same unique equilibrium as the stage-game. Decisions were made simultaneously in each period with no communication among players. And most importantly, assignment of subjects to teams and the matching between teams were randomized for each period. These procedures hinder any effective form of reciprocation among players, rendering tacit collusion impossible or at least unlikely. In the future we intend to examine the more realistic case of a repeated game by keeping the groups' composition constant throughout the interaction. Another interesting line of research involves "mixed" markets consisting of both cooperative and non-cooperative teams.

References

Bertrand, J. (1883) "Review of *Theorie Mathematique de la Richesse Sociale and Recherches sur les Principes Mathematicque de la Theoire des Richesse,*" *Journal des Savants*, 499-508.

Bornstein, G. (1992) "The free rider problem in intergroup conflicts over step-level and continuous public goods," *Journal of Personality and Social Psychology*, 62, 597-606.

Bornstein, G., Budescu, D., & Zamir, S. (1997) "Cooperation in intergroup, two-person, and n-person games of Chicken," *Journal of Conflict Resolution*, 41, 384-406.

Davis, D. D. and Holt, C.A. (1993) *Experimental Economics*, Princeton: Princeton University Press.

Dufwenberg, M. and Gneezy, U. (2000) "Price competition and market concentration: An experimental study," *International Journal of Industrial Organization*, 18, 7-22.

Erev, I., and Roth, A. (1999) "Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria," *American Economic Review*, 88, 848-881.

Fudenberg, D. and Levine, D.K. (1998) *The Theory of Learning in Games*, Cambridge, MA: MIT Press.

Palfrey, T. R. and Rosenthal, H. (1983) "A strategic calculus of voting," *Public Choice*, 41, 7-53.

Rapoport, A. and Bornstein, G. (1987) "Intergroup competition for the provision of binary public goods," *Psychological Review*, 94, 291-299.

Roth, A. E., and Erev, I. (1995) "Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term," *Games and Economic Behavior*, 8, 164-212.

Tirole, J. (1994) *The Theory of Industrial Organization*, MIT Press: Cambridge, Massachusetts.

Thorndike, E. L. (1898) "Animal intelligence: An experimental study of the associative processes in animals," *Psychological Monographs*, 2.

Appendix

Instructions: You are about to participate in a decision-making experiment. During the experiment you will be asked to make a large number of decisions, and so will the other participants. Your own decisions, as well as the decisions of the others, will determine your monetary payoff according to rules that will be explained shortly.

You will be paid in cash at the end of the experiment exactly according to the rules. Please keep quiet throughout the entire experiment and do not communicate in any way with the other participants.

The experiment is computerized. You will make all your decisions by entering the information in the specified locations on the screen. Twelve people participate in this experiment, which includes 100 decision rounds. At the beginning of each round, the 12 participants will be divided randomly into four groups of three persons each, and each group will be paired with another group. The pairing will be done randomly by the computer. You have no way of knowing who belongs to your group and who belongs to the other group.

In each new round, the computer will again divide the participants at random into four groups and each group will be paired at random with another group. At the beginning of a round each of you can demand any number of points between 2 and 25. After all the participants have entered their demands, the computer will sum up the number of points demanded by the three members of your group and will compare it with the total number of points demanded by the three members of the other group.

1. If the total demand made by your group is **lower** than that made by the other group, each member of your group will receive the number of points he or she demanded.

2. If the total demand made by your group is **higher** than that made by the other group, each member of your group will receive nothing (0 points).

3. If the total demand made by your group is **equal** to that made by the other group, each member of both groups will receive half the number of points he or she demanded.

At the end of each round you will receive information concerning (a) the total number of points demanded by your group; (b) the total number of points demanded by the other group; (c) the number of points you earned on that round; and (d) your cumulative earnings up to this point. Then we will move to the next round. Remember that for this new round you will be randomly divided into groups.

At the end of the experiment the computer will count the total number of points you have earned and we will pay you in cash at a rate of 10 points = NIS 1.

The instructions for the cooperative treatment were identical except for the following changes in the payoff rules:

1. If the total demand made by your group is **lower** than that made by the other group, each member of your group will receive one third (1/3) of the group's total demand. In other words, the total number of points demanded by the group will be divided equally among the three group members.

2. If the total demand made by your group is **higher** than that made by the other group, each member of your group will receive nothing (0 points).

3. If the total demand made by your group is **equal** to that made by the other group, each member of both groups will receive one sixth (1/6) of the group's total demand. In other words, the total number of points demanded by the group will be divided by two and then divided equally among the three group members.

Following the reading of the instructions, the participants answered a quiz containing three examples. Each example listed the number of points demanded by each of the six players and the participants were asked to fill in the earning for each player. The experimenter went over the examples and explained the payoff rules until they were fully understood. The examples used in the two treatments were, of course, identical.

Tables:

Table 1: Mean bid, mean winning bid, and mean profit for each 12-person cohort summed across all 100 rounds.

Cohort #	Treatment	Mean Demand	Mean Winning Demand	Mean Payoff (NIS)
1	no-co	10.71	9.18	45.9
2	no-co	8.54	7.77	38.8
3	no-co	8.86	7.42	37.0
4	no-co	8.13	7.50	37.4
5	no-co	7.15	6.25	31.1
6	coop	4.29	3.34	16.6
7	coop	6.72	5.13	25.7
8	coop	5.23	4.16	20.6
9	coop	5.98	4.95	24.8
10	coop	4.92	4.19	20.9

Table 2: Mean demand and mean winning bid per 12-person cohort in the last 10-round block

Cohort #	Treatment	Mean Demand	Mean Winning Bid
1	no-co	11.88	10.07
2	no-co	7.79	7.05
3	no-co	8.96	7.62
4	no-co	4.63	4.31
5	no-co	6.34	5.60
6	coop	3.01	2.41
7	coop	3.34	2.43
8	coop	3.23	2.59
9	coop	3.99	3.23
10	coop	3.55	3.14

Figure 1: Mean bid per round: Experimental resulats

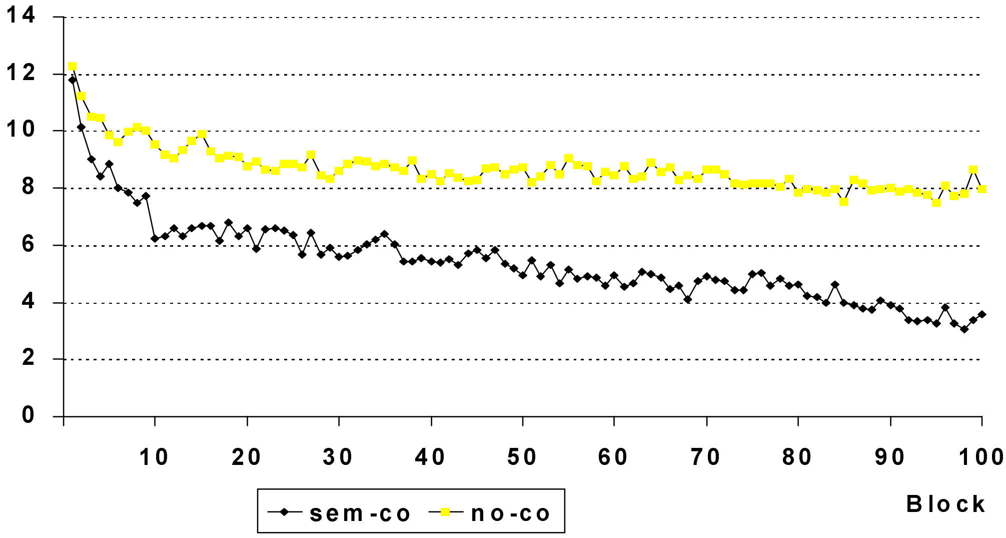

Figure 3: Mean bid per block: Simulated results