Journal of Economic Behavior & Organization
Vol. 66 (2008) 808–821

JOURNAL OF
Economic Behavior
& Organization

www.elsevier.com/locate/econbase

Repeated price competition between individuals and between teams ☆

Gary Bornstein [a,*], Tamar Kugler [b], David V. Budescu [c], Reinhard Selten [d]

[a] *Department of Psychology and Center for Rationality and Interactive Decision Theory, The Hebrew University of Jerusalem, Jerusalem 91905, Israel*
[b] *Department of Management and Policy, University of Arizona, Tucson, AZ 85721, United States*
[c] *Department of Psychology, University of Illinois, 603E Daniel Street, Champaign, IL 61820, United States*
[d] *Department of Economics, University of Bonn, Germany*

Received 18 March 2004; accepted 6 June 2006
Available online 11 January 2007

Abstract

We conducted an experimental study of price competition in a duopolistic market. The market was operationalized as a *repeated* game between two "teams" with one, two, or three players in each team. We found that asking (and winning) prices were significantly higher in competition between individuals than in competition between two- or three-person teams. There were no general effects of team size, but prices increased with time when each team member was paid his or her own asking price and decreased when the team's profits were divided equally. This result is consistent with a simple model of individual learning.

JEL classification: L11; C92

Keywords: Bertrand model; Price competition; Tacit collusion; Experiments; Group behavior

1. Introduction

Oligopolists are typically modeled as unitary, profit-maximizing firms. However, an oligopolist can also be an alliance of firms. Suppose, for example, that a community wants to construct a new

☆ This research was supported by grants from the Israel Science Foundation (907/01 and 535/05) to Gary Bornstein and is part of the EU-TMR ENDEAR Research Network (FMRX-CT98-0238).
* Corresponding author. Tel.: +972 2 5883029; fax: +972 2 5881159.
E-mail address: msgary@mscc.huji.ac.il (G. Bornstein).

0167-2681/$ – see front matter
doi:10.1016/j.jebo.2006.06.004

city hall. The project is too large for any local firm to handle itself; several of them get together to make a common bid. Suppose, further that there are two such composite bidders and that the city will give the contract to the lowest bidder. Assume that it is clear what part of the project each firm in the alliance will complete: one will build the structure, another will provide the landscaping, a third will install the electrical system, and so on. An interesting problem is how much of the profits associated with the common bid should go to each member of the alliance that wins the competition.

Obviously, this is a bargaining problem. The simplest way of modeling the negotiations is to let each firm in the alliance ask for a certain amount, with the common bid being determined as the sum of all these amounts. This admittedly simplistic model has the advantage of being easily amenable to theoretical and experimental investigation. Moreover, even if there is extensive interaction and communication between the business partners, it may be looked upon as cheap talk. In the end, everyone has to make an independent decision as to how much he wants to charge, and the sum of all the demands will be submitted as the group's bid. In this form, the situation can be seen as a non-cooperative game. Of course, communication before the final choice is made would probably have an impact on the final outcome, but in the current investigation we are only studying the purely non-cooperative interaction of making the final demands.

The purpose of this study is to examine how the alliance's size and profit division arrangements affect the outcome of a repeated Bertrand price competition. With this goal in mind, we employ a variation of the original Bertrand game, based on an extension used by Bornstein and Gneezy (2002). The differences between the two games and experimental designs are explicated below.

1.1. The game

Following the seminal work by Fouraker and Siegel (1963), who were the first to study price competition experimentally, most experiments have employed a repeated-game design (see Plott, 1989; Holt, 1995 for reviews of this literature).

The game employed in this study was first studied by Dufwenberg and Gneezy (2000) with individual players. Denote the number of groups in the competition by g, and the number of players in group i by $n(i)$. In our experiment, the game was played by two "groups" with one, two, or three players in each group, that is $g = 2$, and $n(1) = n(2) \in \{1, 2, 3\}$. In each round of the game, each member k in group i simultaneously demands a price $X_{ik} \in \{2, 3, \ldots, 25\}$ (demands are limited to integer numbers). We denote by X_i the sum of demands made by the members of group i ($X_i = \sum_k X_{ik}$). Group i wins the competition if $X_i < X_j$, ties the competition if $X_i = X_j$ and loses it if $X_i > X_j$.

In addition to group size, we manipulated the way in which the profit was divided between the members of the winning group. In the 'private profit' treatment each player is paid her own demand if her group wins and half of her demand in the case of a tie. The payoff of member k in group i is given by

$$\pi_{ik} = \begin{cases} X_{ik} & \text{if} \quad X_i < X_j \\ 0 & \text{if} \quad X_i > X_j \\ \dfrac{X_{ik}}{2} & \text{if} \quad X_i = X_j \end{cases}.$$

In the 'shared profit' treatment, each group member receives an equal share of her group's demand if her group wins and half of that in the case of a tie. The payoff for player k in group i is given by

$$\pi_{ik} = \begin{cases} \dfrac{X_i}{n(i)} & \text{if} \quad X_i < X_j \\ 0 & \text{if} \quad X_i > X_j \\ \dfrac{X_{ik}}{2 \times n(i)} & \text{if} \quad X_i = X_j \end{cases}.$$

Manipulating the teams' profit division arrangement is intended to disentangle the two fundamental problems (free riding and coordination) that distinguish a team from a truly unitary player. When prices are above the competitive equilibrium, the 'private profit' treatment provides each player with an opportunity, indeed a temptation, to free-ride. That is, if the other players in her team settle for a low price, a player can demand a higher price and may yet win. In the 'shared profit' treatment, where equal division of profits is imposed, the opportunity for free riding is eliminated. However, team members still face the problem of coordinating a joint strategy without communicating.[1] Individual players are obviously spared both problems.

Bornstein and Gneezy employed a similar game. They studied groups of three players, and manipulated the profit sharing arrangements (using the terms cooperative and non-cooperative instead of shared profit and private profit, respectively). Our study differs from Bornstein and Gneezy in two important ways. First, we are interested in the effect of group size, in particular in the comparison between truly unitary players (i.e., "groups" of one player only) and groups with multiple players. Second, unlike Bornstein and Gneezy who were mainly interested in individual learning, we focus on the possibility of mutual cooperation or collusion. Therefore, we employ a different matching protocol, a "repeated partners" rather then the "repeated strangers" design used by Bornstein and Gneezy. The implications of these differences are discussed extensively in the next section.

1.2. Theoretical considerations and previous findings

The Nash equilibrium for the stage game, *regardless* of the size of the group or the way profits are divided among group members, is for each player k in either group to demand the minimal price, which in our experiment is $X_{ik} = 2$.

However, when the game is played repeatedly by the same players, as in the present study, the set of equilibria is larger. In an ongoing interaction, behavior can depend on the earlier choices of other players in one's own group and/or the competing group, so outcomes that would be regarded as irrational in a one-shot game may be perfectly rational when the game is repeated. In particular, the collusive outcome, where all players in both teams ask the maximal (i.e., monopoly) price, is supported by Nash equilibrium for a sufficiently large discount factor.

Our objective is not to provide a full analysis of the repeated price-competition game or account for all its strategic aspects. We simply wish to investigate differences in behavior between individuals and teams in a realistic, repeated-game setting. Our primary interest is indeed in the issue of tacit collusion. With repeated interaction, the competitors might be able to collude in a

[1] The 'shared profit' treatment thus amounts to having a committee determine the total bid by aggregating the individual bids submitted and splitting the profit evenly among the players. This arrangement is admittedly not very realistic. Our main interest is clearly in the 'private profit' treatment, where there is a partial conflict of interests within the team, while the 'shared profit' treatment, where internal competition does not exist, serves as a meaningful baseline. Rapoport and Amaldoss (1999) refer to these two distribution rules as "proportional" and "egalitarian", respectively.

purely non-cooperative manner to sustain higher prices than predicted by the one-shot Bertrand model (Tirole, 1988). In the present study, we investigated how the teams' size and their internal profit-sharing arrangements affect the likelihood of collusion.

Team size is predicted to obstruct collusion. As the number of decision makers or players increases, the prospect of successful cooperation typically decreases (e.g., Hamburger et al., 1975). The same is true for coordination (Van Huyck et al., 1990; Ochs, 1995). Hence, the mere complexity of a competition between two- or three-person teams as compared with a competition between individuals renders the realization of common interests in the larger games more difficult. It is possible, however, that the relationship between team size and cooperation is not strictly monotonic. There might be a qualitative difference between individual players, who do not face any internal conflicts and coordination problems, and groups (of any size > 1). The current design allowed us to address this question directly.

Profit-sharing arrangements within the competing groups are also predicted to affect prices in the market. As discussed above, an iterated game makes tacit collusion between the two competitors possible, but if cooperation fails, this setup provides the individual players with the opportunity to learn the structure of the stage-game and adapt their behavior accordingly. Bornstein and Gneezy, who studied price competition between three-person teams of 'repeated strangers', reported that the pace of individual adaptation depends largely on the teams' profit-division arrangement. They found that convergence to the competitive price was much faster in the 'shared profit' treatment than in the 'private profit' one.[2]

Their explanation for this outcome is based on a rather simple process of individual adaptation. In both treatments a high demand by player *k* is likely to result in *k*'s team losing the competition and *k* receiving a payoff of zero. However, if *k*'s team ends up winning the competition, *k*'s payoff is higher in the 'private profit' than in the 'shared profit' treatment, where *k* has to share the profit equally with the other teammates. Since high demands are punished more consistently in the 'shared profit' treatment than in the 'private profit' treatment (where high demand can occasionally lead to high personal profit), players learn to reduce prices faster in the former treatment.

To the extent that this adaptation process plays a role in the current 'repeated-partners' design, we conjecture that prices will remain higher in competition between 'private profit' groups than in competition between 'shared profit' groups. However, the 'partners' design may foster reciprocity, collusion and learning by the specific team members while opponents mask the effects of learning the equilibrium.

2. Experimental procedure

2.1. Subjects and design

The participants were 264 undergraduate students (65 percent females) at the Hebrew University of Jerusalem. Participants were recruited by campus advertisements offering monetary rewards for participating in a decision-making experiment. They had no previous experience with this task. Players participated in the experiment in cohorts of 12. Eight cohorts took part in the

[2] Gunnthorsdottir and Rapoport (2006) found similar results in an inter-group public goods game where the prize for winning the inter-group competition was divided either equally or in proportion to each member's contribution. However, the different structure of their competition leads to theoretically higher contribution levels in the proportional division, while in Bornstein and Gneezy the profit sharing does not alter the theoretical solution.

Table 1
Experimental design and sample sizes

Profit-division arrangement	Team size		
	1	2	3
Shared profit	–	12 obs. (48 subjects)	12 obs. (72 subjects)
Private Profit	–	12 obs. (48 subjects)	12 obs. (72 subjects)
	12 obs. (24 subjects)	–	–

two-person team treatments, and 12 cohorts in the three-person team treatments. Half of these cohorts were in the 'private profit' treatment and half in the 'shared profit' treatment. Finally, two cohorts participated in the individual treatment. Thus, we have 12 independent observations in each of the five cells in our design, as shown in Table 1.

2.2. *Procedure*

Upon arrival at the laboratory each participant received a payment of NIS 10 (approximately, $2.5 at the time of the experiment) for showing up and was seated in a separate cubicle in front of a personal computer. The participants were given written instructions concerning the rules and payoffs of the game (see Appendix B) and were asked to listen carefully while the experimenter read the instructions aloud. Then participants were asked a few questions to test their understanding of the rules. The experimenters checked their answers, and when necessary, the explanations were repeated. Participants were told that to ensure the confidentiality and anonymity of the decisions, they would receive their payments in sealed envelopes and leave the laboratory one at a time with no opportunity to meet the other participants.

The 12 participants were randomly divided into one-, two-, or three-person "teams", depending on the treatment, and each team was randomly matched with another team of the same size for the duration of the experiment. The participants played 100 rounds of the game. The number of rounds was not revealed to the players in advance.

At the beginning of each round each player had to enter a demand of 2–25 points. The computer calculated the total number of points demanded by each team, and the team with the lower total demand won. Each member of that team was paid either her own asking price (in the 'private profit' treatment) or an equal share of the group's profit (in the 'shared profit' treatment). Following each round the participant received feedback concerning (a) the total number of points demanded by the members of his or her team on that round, (b) the total number of points demanded by the members of the competing team on that round, (c) the number of points he or she earned on that round, and (d) his or 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, and the participants were dismissed individually.

3. Results

3.1. *Mean prices are higher in the individual treatment than in the group treatments*

Table 2 presents the mean asking price per player and the mean winning price (summed across all rounds), as well as their mean ranks, for each of the five treatments. These means were analyzed

Table 2
Mean asking price (AP) and mean winning price (WP) of the five treatments across 100 rounds (mean ranks in parenthesis)

Profit-division arrangement		Team size		
		1	2	3
Shared Profit:	AP	–	9.62 (29.50)	7.26 (20.50)
	WP		7.79 (28.92)	5.79 (20.25)
Private Profit	AP	–	(31.58)	8.29 (24.33)
	WP		8.42 (32.17)	7.35 (26.58)
Overall	AP	15.45 (46.58)	9.69	7.78
	WP	13.55 (44.58)	8.11	6.57

using the non-parametric Kruskal–Wallis test. This test ranked the 60 independent observations in our experiment (12 observations × 5 treatments) from the lowest to the highest. The difference in mean ranks among the five treatments is statistically significant for both the mean asking price and the mean winning price ($\chi^2_{(4)} = 15.69$, $p < 0.05$ and $\chi^2_{(4)} = 12.75$, $p < 0.05$, respectively).

This effect is attributed mainly to the fact that prices in the individual treatment were significantly higher than those in the group treatments. We decomposed the effect into four orthogonal contrasts (Marascuilo and McSweeney, 1977). Comparing the individual treatment to the four group treatments reveals a significant difference both for the mean asking price and the mean winning price ($\chi^2_{(1)} = 12.72$, $p < 0.05$ and $\chi^2_{(1)} = 9.75$, $p < 0.05$, respectively). A comparison between the two-member and the three-member group treatments indicates that although prices are somewhat higher in competition between two-member groups, the difference is not statistically significant ($\chi^2_{(1)} = 2.59$, $p > 0.05$. and $\chi^2_{(1)} = 2.00$, $p > 0.05$, for the mean price and the mean winning price, respectively). The differences between the 'shared' and 'private' profit treatments in the mean price and the mean winning price were also not significant ($\chi^2_{(1)} = 0.34$, $p > 0.05$, and $\chi^2_{(1)} = 0.90$, $p > 0.05$, respectively), nor were the interactions between group size and profit division arrangement ($\chi^2_{(1)} = 0.03$, $p > 0.05$, and $\chi^2_{(1)} = 0.09$, $p > 0.05$, respectively). We conclude that prices are generally higher in competition between individuals than in competition between teams. However, neither the size of the teams nor the profit-sharing arrangement within them has an influence on average prices.

3.2. Prices increase over time in competition between individuals and remain stable in competition between teams

Next we turn our attention to the dynamics of prices over time. To facilitate presentation of the results and minimize the effects of trial-to-trial fluctuations, the 100 rounds were aggregated in 10 blocks of 10 consecutive rounds each. Fig. 1a and b, respectively, present the mean asking prices and the mean winning prices per block for each group size.

To analyze the changes in asking prices over time, we computed for each observation the Kendall rank-correlation (τ_b) between the mean asking and winning price per (10-trial) block on the one hand and the block number on the other. The correlation is positive if prices increase over time, negative if they decrease over time, and 0 if there is no systematic trend. Table 3 summarizes the distribution of these correlations and their mean values for each of the five treatments.

We subjected the Kendall correlations to a Kruskal–Wallis test, which indicated an overall treatment effect for the mean asking price ($\chi^2_{(4)} = 10.47$, $p < 0.05$) but not for the mean winning price

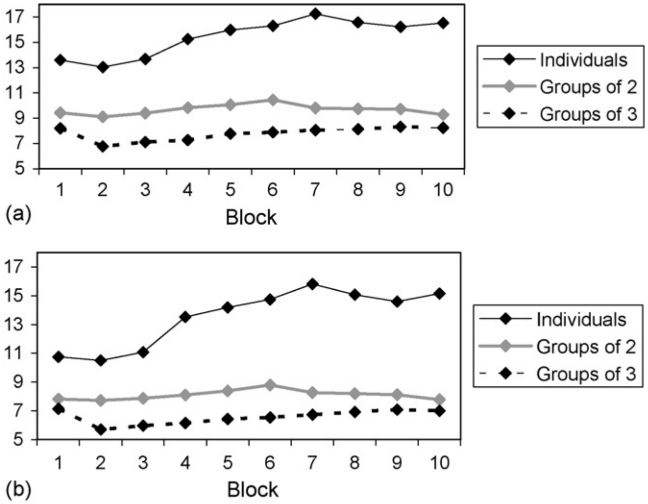

Fig. 1. (a) Mean asking price; (b) mean winning price.

($\chi^2_{(4)} = 8.37, p > 0.05$). When the global test is decomposed into four orthogonal contrasts, we find no significant differences between the individual and the group treatments ($\chi^2_{(1)} = 1.70, p > 0.05$, and $\chi^2_{(1)} = 2.38, p > 0.05$, for the mean asking price and the mean winning price, respectively). The 2-person groups and the 3-person groups are also not significantly different ($\chi^2_{(1)} = 1.96, p > 0.05$, and $\chi^2_{(1)} = 0.52$, $p > 0.05$, for the mean asking price and the mean winning price, respectively). We did, however, find a significant difference between the 'private-profit' and the 'shared-profit' treatments in both the mean asking price and the mean winning price ($\chi^2_{(1)} = 6.80$, $p < 0.05$ and $\chi^2_{(1)} = 5.30$, $p < 0.05$, respectively), but there was no significant interaction between group size and profit-sharing arrangement ($\chi^2_{(1)} = 0.00$, $p > 0.05$, and $\chi^2_{(1)} = 0.17$, $p > 0.05$, for the mean asking price and the mean winning price, respectively).

Table 3
Mean rank correlation (proportion of positive correlations) of block number with mean asking price and mean winning price

Number of players	Profit division arrangement		
	Shared	Private	Mean
1	–	–	(8/12) 0.23* (8/12)
2	−0.28* (3/12) −0.16 (3/12)	0.02 (8/12) 0.05 (8/12)	−0.13 (11/24) −0.06 (11/24)
3	−0.10 (3/12) −0.12 (4/12)	0.19 (9/12) 0.20 (9/12)	0.05 (12/24) 0.04 (13/24)
Mean	−0.19* (6/24) −0.14 (7/24)	0.11 (17/24) 0.13 (17/24)	

Note: *mean correlation significantly different from 0 by a Wilcoxon test ($p < 0.05$).

3.3. Prices increase in competition between 'private-profit' teams and decrease in competition between 'shared-profit' teams

The previous analysis documents a differential trend of prices over time in the two types of profit-sharing arrangements. To investigate further, we performed Mann–Whitney *U*-tests and found the rank-correlations to be significantly higher in the private-profit than in the shared-profit treatment ($p < 0.01$ in a one-sided test when these two group sizes are combined). Similar results were obtained for the winning prices. There was no significant difference between the rank-correlations of individuals and 'private-profit' groups, but there was a significant difference between individuals and 'shared-profit' groups ($p < 0.05$ for asking prices and for winning prices).

Of course, this analysis cannot tell whether the correlations in each treatment are mostly positive or mostly negative (the Mann–Whitney *U*-test would yield identical results whether all correlations are positive or all are negative). Therefore, we counted the number of positive and negative correlations in each profit-sharing treatment across the two group-size treatments (see Table 3). We found that the distribution of positive and negative correlations is distinctly affected by the profit-sharing arrangement. The correlations between asking price and block number are mostly positive (17 out of 24) in the 'private-profit' treatment, indicating that prices often increase over time, and mostly negative (18 out of 24) in the 'shared-profit' treatment, indicating that prices frequently decrease over time. The difference between the two distributions is statistically significant by a Fisher Exact Test ($p < 0.01$). The same is true for the distributions of the winning price.

To summarize, prices tend to increase over time when profits are private and to decrease over time when profits are shared. This trend is depicted in Fig. 2a and b for the mean asking price and the mean winning price, respectively (summed over the two group sizes).

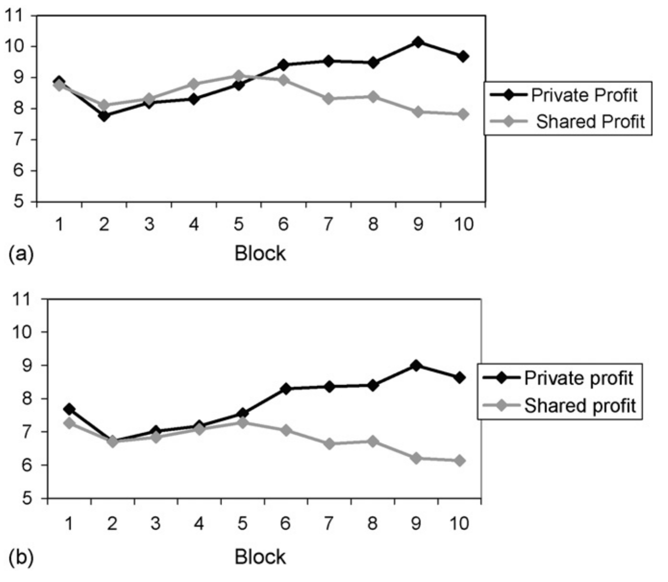

Fig. 2. (a) Mean asking price; (b) mean winning price.

Table 4
Distribution of values of tied games in the five treatments

Type of player	Number of ties	Tied at		
		2 (Nash equilibrium) (percent)	2 < Price < 25 (percent)	25 (Efficient outcome) (percent)
Individual	441	22	21	57
Private profit team of 2	118	1	99	0
Shared profit team of 2	111	17	83	0
Private profit team of 3	115	8	92	0
Shared profit team of 3	72	10	90	0

3.4. Individuals manage to collude more efficiently than teams

Finally, we looked at the occurrences of ties. Ties are observed more often in the individual than the group treatments: 36.75 percent of the games between individuals end up in a tie, compared to only 8.67 percent in the multiple-player groups.[3] Moreover, the rate of ties in the individual treatment increased systematically as the game progressed (from 22.5 percent in the first block to 46.67 percent in the last one), whereas, in the other treatments it remained quite stable across blocks. In Table 4, we distinguish between three types of ties: (a) all players request 2, the competitive price, which is the single Nash equilibrium of the stage game, (b) all players request 25, the highest possible price, which is the joint profit maximizing outcome, and (c) the average request per player is some other amount between 2 and 25.

There is a significant difference ($\chi^2_{(8)} = 467.3$; $p < 0.05$) between the patterns of ties recorded in the single- and multiple-player groups; most ties (57 percent) between individuals are joint-profit maximizing (i.e., monopoly) prices, and the number of efficient ties increases from 11.67 percent in block 1–30 percent in block 10, but there is not even one instance of an efficient tie between multiple-player groups! Interestingly, the average payoff per player in tied games is inversely related to group size: 17.37 for single players, 8.02 for dyads and 5.87 for triads. Examining the pattern of ties in the multiple-player groups reveals a higher fraction of ties that are competitive Nash equilibria in the 'shared-profit' treatments (26/183 = 14.2 percent) than in the 'private-profit' treatments (10/233 = 4.3 percent). This difference is significant (by a test of equality of proportions: $Z = 3.48$; $p < 0.05$), indicating again that 'shared-profit' teams are more competitive than are 'private-profit' ones.

3.5. Individual learning

To investigate how individual learning of the game structure affects the observed dynamics we employed the reinforcement learning model formulated by Roth and Erev (1995) and Erev and Roth (1998). The basic principle underlying this model is Thorndike's "Law of Effect" (Thorndike, 1898), which states that choices that have led to good outcomes in the past are more likely to be

[3] The 2400 *individual* asking prices can be used to calculate the probability of tied games in every condition by chance alone, assuming that there is no difference over the 100 rounds and that the players are independent (both within and between groups). The probability of a chance tie is 0.111 for single players, 0.033 for dyads, and 0.021 for triples. Clearly, there are more ties than expected by chance in all cases, indicating some tacit coordination among players, but the rate is considerably higher for individuals than for groups.

Table 5
Correlations of changes in individual bids with outcomes and payoffs on the previous round in the five treatments

	Individuals	Groups of 2	Groups of 3
Outcomes			
Private profit		Mean = 0.27; S.D. = 0.17	Mean = 0.25; S.D. = 0.19
Shared profit		Mean = 0.29; S.D. = 0.18	Mean = 0.27; S.D. = 0.16
	Mean = 0.36; S.D. = 0.16		
Payoffs			
Private profit		Mean = 0.12; S.D. = 0.17	Mean = 0.12; S.D. = 0.18
Shared profit		Mean = 0.15; S.D. = 0.19	Mean = 0.18; S.D. = 0.16
	Mean = 0.12; S.D. = 0.21		

repeated in the future. Before fitting the model, we illustrate this basic empirical regularity by correlating the change in individual bids (round t bid − round $(t-1)$ bid) with results of the previous $(t-1)$ round (1 = win, 0 = tie, −1 = lose) and payoffs on the previous round $(t-1)$. The means and standard deviations of the individual correlations are presented in the two panels of Table 5.

The mean correlations are positive for all conditions, indicating that players increase their bids following a win and decrease their bids after a loss, and that these changes are positively associated with the actual payoffs. Kruskall–Wallis tests found no significant differences between the five groups in this respect ($\chi^2_{(4)} = 7.45$, $p > 0.05$ for outcomes, and $\chi^2_{(4)} = 5.10$, $p > 0.05$ for payoffs).

In applying the Roth and Erev model, we assumed that Player k considers $S = 24$ pure strategies (choosing an integer between 2 and 25). At time $t = 1$, the player's initial propensities to choose each of these strategies are equal. Her propensity to select a particular strategy, say strategy s, at time $t = 1$ is denoted by $q_{ks}(1)$. If Player k plays strategy s at time t and receives a payoff of x, her propensity to play s is updated by setting $q_{ks}(t+1) = q_{ks}(t) + x$, while for all other pure strategies, propensities do not change. The probability $p_{ks}(t)$ that Player k will play strategy s at time t is $p_{ks}(t) = q_{ks}(t)/\Sigma q_{ks}(t)$, i.e., the ratio of the propensity to play the sth 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 of the propensities to play strategies s and not-s at time $t = 1$, and the sum of the initial propensities over the 24 strategies, $S(1) = q_{ks}(1) + q_{ks'}(1)$. Following Roth and Erev, we set $S(1) = 10$.

The predictions of the learning model were derived by having simulated subjects interact with each other in the experimental environments. Fig. 3 presents the mean bids in 400 simulated groups, each playing 100 rounds of the game under either the 'private' or the 'shared' profit treatments. The top panel displays the results as a function of the group size and should be compared with Fig. 1a. The bottom panel, which presents the results of the groups by profit-division arrangement, should be compared to Fig. 2b. Some predictions of the learning model are visibly out of line with our experimental findings. In particular, the model forecasts that prices should decline faster in competition between individuals than between groups, while in the experiment prices actually increased in the individual treatment and did not change much during the course of the game in the group treatments. Other predictions are supported by the experimental data. Most notably, the model predicts that prices in the 'shared' profit treatment should decline with time while those in the 'private' profit treatment would remain quite stable, and we indeed observed an increased price gap between the two profit division arrangements conditions over time.

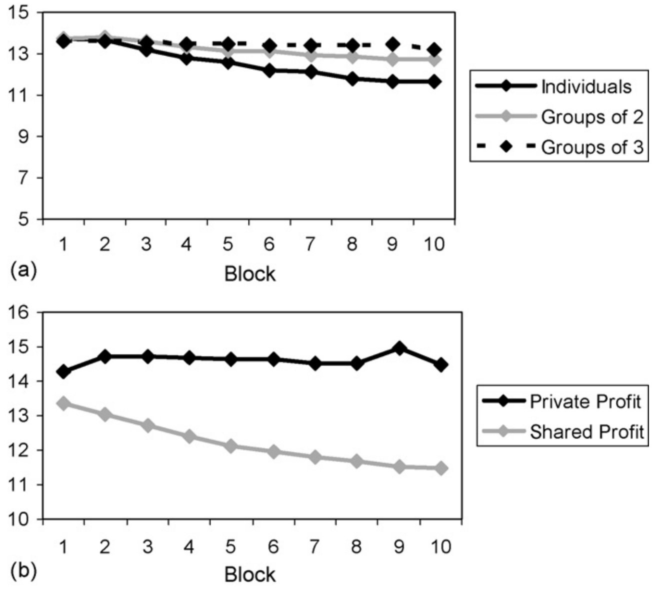

Fig. 3. (a) Mean asking price—simulation; (b) mean asking price—simulation.

3.6. Comparison to Bornstein and Gneezy

Fig. 4 presents the mean asking prices reported by Bornstein and Gneezy. Because the games used identical parameters, these results can be compared directly to Fig. 2a (Bornstein and Gneezy used groups of three players only, while Fig. 2a exhibits the combined results from groups of two and three members, but see Section 3.1).

There are some very distinct differences between the results of the two experiments. First, in Bornstein and Gneezy the difference between the two profit sharing conditions is evident after as little as three rounds of the game, while in our experiment, the mean asking price remains very stable for the first 50 rounds. Second, while in Bornstein and Gneezy prices decline in both conditions (albeit much faster in the cooperative, profit sharing condition), in our experiment prices drop only in the shared profit condition and actually increase in the private profit condition.

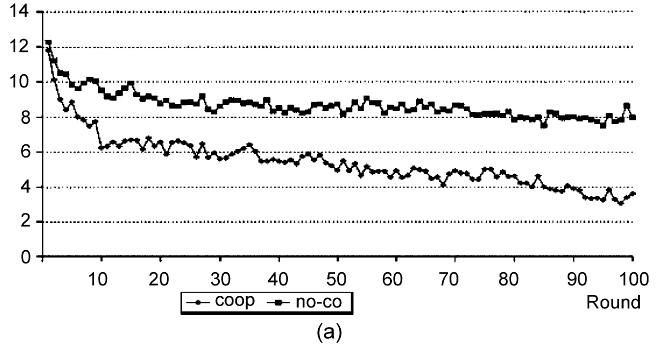

Fig. 4. Mean asking price per round—Bornstein and Gneezy (2002).

The most plausible explanation for these differences is the matching protocol. While the repeated strangers design (Bornstein and Gneezy) only enables players to learn the game structure, the present repeated partners design also allows for other dynamics, such as signaling, reciprocity and collusion to take place simultaneously. Therefore, players in our experiment (especially in the individual condition and the private profit conditions) manage to avoid the cut-throat price competition that causes behavior to converge to the stage game equilibrium prices. This can also explain why the reinforcement learning model, that provided an almost perfect fit to the results of Bornstein and Gneezy, fails to capture some of our main results: this model focuses on payoff reinforcement and does not account for the dynamics of repeated interaction between the same players.

4. Discussion

In the Bertrand game, if firms meet only once and quote their price simultaneously and independently (i.e., non-cooperatively), the prices should theoretically equal the marginal cost, even if there are only two firms in the market. In practice, however, firms often interact repeatedly, which may upset the Bertrand outcome (Tirole). In such repeated interactions, firms must take into account not only current profits but also the potential long-term losses of a price war. These long-term considerations decrease the temptation to cut prices and may encourage the competitors to collude in a purely non-cooperative manner to sustain higher prices than predicted by the one-shot model (Tirole). In fact, Chamberlin (1929) suggested that when the number of firms in the market is small, tacit collusion resulting in the monopoly price is the most likely outcome.

This prediction is based on the simplifying assumption that the competitors operating in the marketplace are unitary players. In reality, however, the competitors often consist of multiple players. In these cases, the competitors' profit division arrangement, and in particular, the possibility of conflicting interests and coordination problems, must be taken into account (Bornstein, 1992, 2003; Rapoport and Bornstein, 1987). This is obviously true when the competitors are alliances of firms (Amaldoss et al., 2000), but it is also true when the competitors are single firms (Kirstein and Kirstein, 2004). Principal-agent theory acknowledges the existence of conflicting interests within firms, but when firms are studied in strategic contexts of competition against other firms, they are typically modeled as unitary players. This is also reflected in experimental markets, where firms are commonly represented by individual subjects (Holt).

The goal of the present study was to investigate whether the market is sensitive to the violation of the unitary player assumption. Toward this goal, we modeled the competitors in a duopolistic market as either individuals or teams. We also varied the profit-sharing arrangements of the competing teams, so that in one treatment each team member was paid his or her asking price, while in the other treatment the team's profit was divided equally among its members.

The fixed-matching design of our experiment rendered tacit collusion between the two competitors both theoretically possible and practically viable. Nonetheless, we found that individual players were much better able to collude than teams. Individuals managed to keep the average winning price above 13 points (out of the maximum of 25), as compared with an average winning price of about 8 points in markets consisting of two-person teams and 6.5 points in competitions between three-person teams. Moreover, in competitions between individuals prices increased with practice, and toward the end of the game the collusive outcome was achieved in a substantial number of cases, whereas, in competition between teams prices remained stable, and there was little evidence of learning to collude. Clearly, this market is highly sensitive to violations of the unitary player assumption, and full collusion at the highest possible price is much less likely when the competitors are multi-player teams.[4] This is obviously good news for consumers. Collusion resulting in high prices is typically considered undesirable, as reflected in antitrust policies.

We also found that profit-sharing arrangements within the competing teams had some effect on the market. While the mean prices in the 'shared' and 'private' profit treatments were not statistically different, the two treatments differed significantly in their dynamics of price change. From about round 50 on, prices tended to increase over time in the 'private profit' treatment and decrease over time in the 'shared profit' treatment. These results are in some respects different from those reported by Bornstein and Gneezy. They found that in a mixed-matching (strangers) design, prices declined in *both* the 'private' and 'shared' profit treatments, but the decline was much faster when profits were shared. These results can be explained by the Roth and Erev reinforcement learning model. This model cannot explain equally well the results obtained in the current fixed-matching (partners). Whereas, individual learning can explain the decrease in prices in the 'shared profit' treatment, it cannot explain the increase in prices in the 'private profit' treatment. We can only surmise that, given the relatively weak incentive to reduce prices in this treatment (as demonstrated by the simulation), enabled the players to avoid the inefficient outcome of a full scale price war, without leading to the monopoly price.

Both studies are in agreement with respect to profit sharing: prices are sustained at a higher level when each team member is paid his or her own asking price than when the team's profits are divided equally. Thus, unlike competition between agents, which lowers prices (Dufwenberg and Gneezy, 2000), competition *within* agents seem to cause an increase in prices over time (at least in this type of market). A similar argument was made recently by Kirstein and Kirstein (2004), in a different context. Modeling the intra-firm conflicts in oligopoly market as a principal-agent problem, they demonstrated that principals can stabilize a cartel by providing inefficient intra-firm incentives.

The present study focused on the symmetric situation where the competing agents were either individuals or teams of equal size, with identical profit-sharing arrangements. It might be interesting to study asymmetric or "mixed" markets consisting of individuals and teams with different profit-sharing arrangements to find out which type of agent (if any) has a decisive effect on the market's behavior. It would also be of much interest to conduct experiments on inter-group price competition that include a stage of cheap talk stage within teams.

Acknowledgement

We thank Simon Gächter, Uri Gneezy and Heike Hennig-Schmidt for their helpful comments.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.jebo.2006.06.004.

[4] This result may not generalize to other types of markets. In a recent study Raab and Schipper (2004) compared the behavior of individuals and teams (with various profit division arrangements) in a quantity (i.e., Cournot) competition and found that the two types of players behaved in the same way.

References

Amaldoss, W., Meyer, R., Raju, J., Rapoport, A., 2000. Collaborating to compete. Marketing Science 19, 105–126.

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., 2003. Intergroup conflict: individual, group, and collective Interests. Personality and Social Psychology Review 7, 129–145.

Bornstein, G., Gneezy, U., 2002. Price competition between teams. Experimental Economics 5, 29–38.

Chamberlin, E., 1929. Duopoly: values where sellers are few. Quarterly Journal of Economics 43, 63–100.

Dufwenberg, M., Gneezy, U., 2000. Price competition and market concentration: an experimental study. International Journal of Industrial Organization 18, 7–22.

Erev, I., Roth, A.E., 1998. Predicting how people play games: reinforcement learning in experimental games with unique, mixed strategy equilibria. The American Economic Review 88, 848–881.

Fouraker, L., Siegel, S., 1963. Bargaining Behavior. McGraw-Hill, New York.

Gunnthorsdottir, A., Rapoport, A., 2006. Embedding social dilemmas in intergroup competition reduces free-riding. Organizational Behavior and Human Decision Processes 101, 184–199.

Hamburger, H., Guyer, M.S., Fox, J., 1975. Group size and cooperation. Journal of Conflict Resolution 19, 503–531.

Holt, C.A., 1995. Industrial organization: a survey of laboratory research. In: Kagel, J.H., Roth, A.E. (Eds.), The Handbook of Experimental Economics. Princeton University Press, Princeton, pp. 349–443.

Kirstein, R., Kirstein, A., 2004. Inefficient Incentives within Firms may Stabilize Cartels in Cournot Oligopolies. University of Bonn, Mimeo.

Marascuilo, L.A., McSweeney, M., 1977. Nonparametric and Distribution-free Methods for the Social Sciences. Brooks/Cole Publishing Company, Monterey, CA.

Ochs, J., 1995. Coordination problems. In: Kagel, J.H., Roth, A.E. (Eds.), The Handbook of Experimental Economics. Princeton University Press, Princeton, pp. 195–249.

Plott, C.R., 1989. An updated review of industrial organization: applications of experimental methods. In: Schmalensee, R., Willig, R.D. (Eds.), Handbook of Industrial Organization. Elsevier, North Holland, pp. 1109–1176.

Raab, P., Schipper, B.C., 2004. Cournot Competition between Teams: An Experimental Study. University of Bonn, Mimeo.

Rapoport, A., Amaldoss, W., 1999. Social dilemmas embedded in between-group competitions: effects of contest and distribution rules. In: Foddy, M., Smithson, M., Schneider, S., Hogg, M. (Eds.), Resolving Social Dilemmas. Psychology Press, Philadelphia, pp. 67–84.

Rapoport, A., Bornstein, G., 1987. Intergroup competition for the provision of binary public goods. Psychological Review 94, 291–299.

Roth, A.E., Erev, I., 1995. Learning in extensive form games: experimental data and simple dynamic models in the intermediate term. Games and Economics Behavior 8, 164–212.

Thorndike, E.L.,1898. Animal Intelligence: An Experimental Study of the Associative Processes in Animals. Psychological Review, Monograph Supplements 8, Macmillan, New York.

Tirole, J., 1988. The Theory of Industrial Organization. MIT Press, Cambridge, Massachusetts.

Van Huyck, J., Battalio, R., Beil, R., 1990. Tacit coordination games, strategic uncertainty, and coordination failure. American Economic Review 80, 234–248.