Bonn Econ Discussion Papers

Discussion Paper 1/2002

Stability of the Cournot Process - Experimental Evidence

by

Steffen Huck, Hans-Theo Normann,
Jörg Oechssler

January 2002

Bonn Graduate School of Economics
Department of Economics
University of Bonn
Adenauerallee 24 - 42
D-53113 Bonn

The Bonn Graduate School of Economics is
sponsored by the

Deutsche Post World Net

MAIL EXPRESS LOGISTICS FINANCE

Stability of the Cournot Process – Experimental Evidence*

Steffen Huck
Department of Economics & ELSE
University College London

Hans–Theo Normann†
Department of Economics
Royal Holloway, London

Jörg Oechssler
Department of Economics
University of Bonn

January 21, 2002

Abstract

We report results of experiments designed to test the predictions of the best reply process. In a Cournot oligopoly with four firms, the best reply process should theoretically explode if demand and cost functions are linear. We find, however, no experimental evidence of such instability. Moreover, we find no differences between a market which theoretically should not converge to Nash equilibrium and one which should converge because of inertia. We investigate the power of several learning dynamics to explain this unpredicted stability.

JEL– classification numbers: C72, C92, L13.

Key words: best reply process, Cournot oligopoly, learning, experiments, imitation.

---

*We thank two referees and participants at the ESA meeting in Mannheim for valuable comments. Financial support by the DFG, grant OE-198/1/1 is gratefully acknowledged.

†Department of Economics, Royal Holloway, University of London, Egham, Surrey, TW20, 0EX, U.K.

1 Introduction

It is well known that the adjustment process suggested by Cournot (1838) for an oligopoly, namely that each firm plays a best reply to the other firms' previous output, is not stable in the general case. Theocharis (1960) shows that in oligopolies with more than two firms and with linear demand and cost functions the best reply process does not converge: if there are three firms, finite oscillations around the equilibrium positions occur, and with four or more firms the process shows explosive fluctuations.

Subsequently, it has been shown that this result is not robust to small changes in the assumptions. In particular, the system can become stable if adjustment to the best reply is only partial (McManus and Quandt, 1961), or if marginal cost are increasing (Fisher, 1961).

In this experiment we test whether convergence to the equilibrium really depends on such intricate details of the model. We propose two treatments for a four–firm oligopoly with linear demand and cost functions which allow to test for the stability of the Cournot adjustment process. Only in treatment A is instantaneous and perfect adjustment possible. In treatment B, firms must stick to last period's quantity with a probability of 1/3. In Huck, Normann, and Oechssler (1999) we prove that such a system with inertia is stable. Thus, the predictions are clear. In the first case, the process should oscillate perpetually between two extreme values. In the second case, the process should converge to equilibrium.

We find, however, no noticeable difference between the two treatments. In both cases average quantities are slightly above, but still rather close to the Cournot equilibrium outputs. The modal choice of individual quantities is at the Cournot outcome in both treatments. In experimental oligopoly markets, Theocharis' instability result does not occur.

There may be several reasons for this (unpredicted) stability. One reason could be that subjects follow a process that smooths best replies, e.g. like fictitious play. Another reason may be that subjects tend to imitate what others do. Social psychologists (e.g. Asch, 1952) have shown a long time ago that imitation is an important factor for explaining learning behavior.[1] We find evidence that imitation of the average quantity of other firms plays an important role in our experiment. Fictitious play, on the other hand, does not explain the data particularly well. Our regression results suggest that,

[1]Recently, imitation (of last period's most successful action) has also been studied in the context of oligopoly by Vega–Redondo (1997).

in the aggregate, subjects mix between playing best replies and imitating others.

These results may be compared to an interesting experiment by Cox and Walker (1998). They also analyze convergence of play in two treatments of which only one is theoretically stable. In Cournot duopoly the best-reply dynamic is stable only if firm 1's reaction function is steeper than firm 2's in the neighborhood of the equilibrium (with firm 1's quantity on the horizontal axis). Cox and Walker show that there is a sharp distinction between these two cases supporting the theoretical results. Play is almost never near the theoretically unstable equilibrium but converges nicely to the theoretically stable equilibrium.

Another related paper is by Rassenti et al. (2000) who ran several five-firm oligopoly experiments. Their central issue is whether repeated play will yield convergence to the unique and theoretically unstable static non-cooperative Nash equilibrium. Surprisingly, and in contrast to previous studies and our own, they observed convergence only at the aggregate level. There is no convergence at the individual level. A possible explanation for this is that Rassenti et al. introduce substantial asymmetries in cost. A further difference to our paper is that they do not have a control treatment which is theoretically stable.

The remainder of this paper is organized as follows: Section 2 introduces the experimental design. Section 3 offers theoretical predictions, while Section 4 presents the experimental results. In Section 5 we conclude.

2 Experimental design

In a series of computerized[2] experiments we studied a homogeneous multi–period Cournot market with linear demand and cost. There were four symmetric firms in each market. Quantities could be chosen from a finite grid between 0 and 100 with .01 as the smallest step. The demand side of the market was modelled with the computer buying all supplied units according to the inverse demand function

$$p^t = \max\{100 - Q^t, 0\}, \tag{1}$$

[2]We are grateful to Klaus Abbink and Karim Sadrieh for providing us with their software toolbox "RatImage" (Abbink and Sadrieh, 1995) which we used for the programming of the experiments.

with $Q^t = \sum q_i^t$ denoting total quantity in period t. The cost function for each seller was simply

$$C(q_i^t) = q_i^t. \tag{2}$$

Hence, profits were

$$\pi_i^t = (p^t - 1)q_i^t. \tag{3}$$

The unique Cournot Nash equilibrium of the stage game is given by

$$q_i^N = \frac{100 - 1}{5} = 19.8,\ i \in I, \tag{4}$$

yielding a price of $p^N = 20.8$. The collusive outcome would be at $q_i^C = 12.375$ resulting in a price of $p^C = 50.5$.

The number of periods was 40 in all sessions and this was commonly known. Subjects possessed all essential information about the market, i.e. they were informed about the symmetric demand and cost functions in plain words.[3] Furthermore, subjects had the possibility to use a 'profit calculator', which served two functions. A subject could enter some arbitrary 'total quantity of other firms'. Then he could either enter some amount as his own quantity in which case the calculator informed him about the resulting price and his resulting personal profit. Or, he could press a 'Max'–button in which case he was informed about the quantity which would yield him the highest payoff given the total amount of others. Additionally, the calculator computed price and profit for this best response.[4] This function was designed to give the best-reply process the best chance possible. The calculator was used, on average, in two of out three periods.

After each market period subjects were informed about the total quantity the others had actually supplied, about the resulting price and their personal profits. Additionally, they were reminded of their own quantity. When deciding in the next period this information remained present on the screen. Results of earlier periods were, however, not available, but subjects were allowed to take notes and a few did.

There were two treatments. In treatment A subjects could adjust their quantities in every period. In treatment B we introduced some inertia: After round one, chance moves, which were independent across individuals, determined in each period whether a subject was allowed to revise his quantity decision. This was done by a 'one–armed bandit' which appeared on

[3] Since we recruited many non-economics students as subjects, we were careful not to use any formulas or technical terms in the instructions.

[4] In the experiments we did not use the expression 'best response.'

the screen showing three equiprobable numbers "0","1", and "2". If "0" occurred no adjustment was allowed. Hence, the probability for allowing revision was 2/3.

The experiments were conducted in April and May 1997 in the computer lab of the economics department of Humboldt University. All subjects were recruited via posters from all over the campus. Almost half of the subjects studied fields other than economics or business and had no training in economics at all.

In each session eight subjects participated, constituting two groups of four firms. Subjects were randomly allocated to computer terminals in the lab such that they could not infer with whom they would interact in a group of four. For both treatments we had six groups of subjects — making a total of 48 subjects who participated in the experiments.

Subjects were paid according to their total profits. Profits as in (3) where denominated in 'Taler', the exchange rate for German Marks (500:1) was known. Since we considered the Theocharis result as a possible outcome in treatment A, we wanted to make sure — besides the usual bankruptcy problems — that subjects would not be frustrated by low or negative payoffs.[5] So, additionally subjects earned a fixed payoff of Taler 150 each round. The average payoff was about DM 37.84 which at the time were roughly $21. Experiments lasted 60 minutes including instruction time.

Instructions (see Appendix A) were written on paper and distributed in the beginning of each session. After the instructions were read we conducted one trial round in which the different windows of the computer screen (see Figure 1 in Appendix B) were introduced and could be practiced. When subjects were familiar with both, the rules and the handling of the computer program, we started the first round.

3 Theoretical predictions

In this section we analyze the implications of several learning dynamics — best reply (BR), fictitious play (FP), imitation of average behavior (AV), and a mixed process. For each theory we analyze whether the process converges and, if so, where it converges to. Furthermore, we calculate for each process theoretical autocorrelations, which can then be compared to the empirically

[5] See Holt (1985, p. 317) for the argument that the usual promises in the instructions that one can earn a "considerable amount of money" might bias subjects against zero-profit outcomes.

observed ones.

3.1 Best-reply processes

The basic best-reply process, suggested by Cournot, assumes that each firm plays a best reply to the other firms' output from last period. If all players use this rule, i.e. if for all i, $q_i^t = \frac{1}{2}\left(99 - Q_{-i}^{t-1}\right)$, we get the following process for total quantities

$$Q^t = 198 - \frac{3}{2}Q^{t-1}. \tag{5}$$

Clearly, the system explodes. Due to the non–negativity constraint for quantities, individual quantities eventually oscillate between zero output and the monopoly output. Total quantities oscillate between zero and four times the monopoly output, i.e. $Q^t = 198$.

Consider now a best reply process with inertia and let θ denote the probability that a firm must stick to its previous quantity. The inertia stabilizes the process. In Huck, Normann, and Oechssler (1999, Prop. 1) we show formally that the resulting Markov process converges globally to the Cournot equilibrium for any $\theta \in (0, 1)$. The proof is based on the theory of potential games (Monderer and Shapley, 1996) and proceeds by constructing an improvement path from any arbitrary state to the equilibrium. The equilibrium is, of course, an absorbing state with respect to best-reply dynamics.

While the process with inertia is stochastic, we can calculate an average autocorrelation coefficient by noting that (on average)

$$q_i^t = \theta q_i^{t-1} + (1-\theta)\left(\frac{99 - Q_{-i}^{t-1}}{2}\right) \tag{6}$$

and so

$$Q^t = 198(1-\theta) - \frac{3-5\theta}{2}Q^{t-1}. \tag{7}$$

Thus, there should be negative autocorrelation as long as $\theta < \frac{3}{5}$. In treatment B we have $\theta = \frac{1}{3}$. Hence, theory predicts a negative autocorrelation of $-\frac{2}{3}$ if all subjects play best replies.

An alternative way of smoothing a best reply process is through fictitious play (FP) beliefs (Robinson, 1951). In its basic version players choose in each round a best reply against the relative frequency of the combined quantities Q_{-i}^{τ} of the remaining firms in periods $\tau = 1, ..., t-1$. Given the linear structure of the Cournot oligopoly, this implies that player i chooses a best reply against $\frac{1}{t-1}\sum_{\tau=1}^{t-1} Q_{-i}^{\tau}$ (after choosing an arbitrary strategy in round

1). In particular, for $t > 1$

$$q_i^t = \frac{1}{2}\left(99 - \frac{\sum_{\tau=1}^{t-1} Q_{-i}^{\tau}}{t-1}\right). \tag{8}$$

It is well known (see Monderer and Shapley, 1996) that fictitious play converges to a Nash equilibrium in potential games, i.e. in our case to the unique Cournot equilibrium. With inertia parameter θ, we have for total quantities

$$Q^t = 198(1-\theta) + \theta Q^{t-1} - \frac{3}{2}(1-\theta)\frac{\sum_{\tau=1}^{t-1} Q_{-i}^{\tau}}{t-1}. \tag{9}$$

Hence, for $\theta = 1/3$ an autocorrelation of $\frac{1}{3} - \frac{1}{t-1}$ results. Without inertia ($\theta = 0$) the process yields an autocorrelation of $-\frac{3}{2(t-1)}$.

3.2 Imitation

Research in psychology and social biology shows that individuals often "learn" by imitating others especially in complex environments. In a famous experiment Asch (1952) found that people who are faced with similar decisions tend to follow the decisions of other members in their group. In our experimental setting subjects could not observe individual quantities. But they were able to observe total — and therefore average — quantities of the remaining firms. It seems reasonable that subjects who are uncertain about what to do and observe that the average quantity of the other firms deviates from their own quantity, imitate this average quantity — thinking along the line of "everyone else can't be wrong". A preference for cautious behavior and a taste for conformity could be further reasons for imitating the average.

This would result in the following process:

$$q_i^t = \frac{Q_{-i}^{t-1}}{3}. \tag{10}$$

If all subjects were to follow this rule, clearly the process is bounded above and below by the highest and lowest initial quantities. Without inertia the process would converge simply to the average of all starting values, as can be seen by solving the system of equations (10) recursively, which yields

$$q_i^t = \frac{Q^1}{4} + \frac{3q_i^1 - Q_{-i}^1}{4} \cdot \left(-\frac{1}{3}\right)^{t-1}. \tag{11}$$

With inertia the process depends on the realizations of the randomization device and is therefore path dependent. Nevertheless, on average quantities are given by

$$q_1^t = \theta q_1^{t-1} + (1-\theta)\frac{Q_{-i}^{t-1}}{3} \tag{12}$$

which yields $Q^t = Q^{t-1}$ independently of θ. Thus, total quantities should be constant on average and the autocorrelation coefficient should be 1.

3.3 Mixed process

In anticipation of our experimental results we derive here some properties of a mixture of best reply and "imitate the average". Apart from the fact that such a mixed process seems to suggest itself given our data, it could result because all subjects actually mix, or – more plausibly – because some subjects play best reply and others imitate (see Gale and Rosenthal, 1999, for a justification of a similar mixed process).[6] Generally, if α denotes the weight given to best replies and $1-\alpha$ the weight given to the average quantity of the other subjects, we get the following difference equation for the case without inertia

$$q_i^t = \alpha\frac{99 - Q_{-i}^{t-1}}{2} + (1-\alpha)\frac{Q_{-i}^{t-1}}{3}. \tag{13}$$

Thus total quantities follow the process

$$Q^t = 198\alpha + \frac{2-5\alpha}{2}Q^{t-1}. \tag{14}$$

This difference equation is stable if $\alpha < \frac{4}{5}$. If it is stable, it converges to the Cournot equilibrium. Autocorrelation is positive for $\alpha < \frac{2}{5}$ and negative otherwise.

With inertia the mixed process yields (on average)

$$q_i^t = \theta q_i^t + (1-\theta)\left[\alpha\frac{99 - Q_{-i}^{t-1}}{2} + (1-\alpha)\frac{Q_{-i}^{t-1}}{3}\right]. \tag{15}$$

Summing over i gives

$$Q^t = 198\alpha(1-\theta) + \left(1 - \frac{5}{2}\alpha(1-\theta)\right)Q^{t-1}. \tag{16}$$

Thus, for $\theta = 1/3$ convergence is assured independently of α, and average autocorrelation is given by $1-\frac{5}{3}\alpha$. We summarize the convergence properties and theoretical autocorrelations in Table 1.

[6] They consider a population in which most people imitate the average but some experiment on a trial & error basis. Since the presence of experimenters has similar effects as that of best-reply players, it is not surprising that they find convergence to the Nash equilibrium and stability in the large. However, due to the stochastic nature of experimentation some interesting instabilities arise in the small.

Table 1: Theoretical predictions

Treatment	Process	Convergence	Autocorr.
A	best reply	no	–
	FP	yes	$-\frac{3}{2(t-1)}$
	imitation	yes	1
	mix	if $\alpha < 4/5$	$1-\frac{5}{2}\alpha$
B	best reply	yes	$-2/3$
	FP	yes	$\frac{1}{3}-\frac{1}{t-1}$
	imitation	yes	1
	mix	yes	$1-\frac{5}{3}\alpha$

Note: α denotes the weight given to best reply.

Table 2: Summary Statistics

Treatment	Mean35	Mean20	Avg σ_{0-20}	Avg σ_{20-40}	Autocorr.
A	82.83 (6.17)	83.98 (6.79)	13.57	8.15	−.043
B	84.32 (4.56)	82.56 (2.48)	16.41	10.00	.387

Note: Mean35 (Mean20) is the average total quantity measured over the last 35 (20) periods and over all six groups in one treatment. Standard deviations in parentheses. Avg $\sigma_{t-t'}$ denotes the average standard deviation of total quantities from round t to t'.

4 Experimental results

Table 2 reports average total quantities for the last 35 and the last 20 periods, respectively.[7] The mean quantity in both treatments is only slightly above the Cournot–Nash quantity of 79.2. Standard non–parametric tests show that there are no significant differences between the mean quantities for treatment A and B at any reasonable significance level.

Also shown are average standard deviations of total quantities over the first and the second half of the experiment.[8] It can be seen that the standard deviations of total outputs are considerably smaller in the second half of the experiment. This is significant for both treatments ($p = 5.8\%$ in treatment A, $p = 1.4\%$ in treatment B (Wilcoxon test)) and indicates convergence. Clearly, there is no tendency for explosive behavior in treatment A. Some-

[7] The complete data set is available at http://www.wiwi.uni-bonn.de/with2/oechssler

[8] For each group of firms we calculated the standard deviations of total quantities over time. Table 2 reports the treatment averages of these standard deviations. For the tests below each group counted as one observation.

Table 3: Hit ratios

Treatment		$z < 0$	$0 \leq z < .8$	$0.8 \leq z < 1$	$z = 1$	$1 < z \leq 1.2$	$z > 1.2$
A	BR	21.5	52.0	4.6	5.2	4.1	12.6
	FP	49.7	22.5	4.2	0	4.7	18.8
	AV	35.3	46.3	2.5	0.5	2.4	13.0
B	BR	26.4	43.0	5.4	5.4	3.1	16.6
	FP	46.0	22.1	5.3	0	4.3	22.3
	AV	35.7	44.6	4.1	1.0	2.3	12.3

Note: Only rounds in which subjects were allowed to adjust their quantities are included.

what surprisingly, average standard deviations are even lower in treatment A, although, this difference is not significant. Finally, Table 2 also reports empirical autocorrelation coefficients which are analyzed below.[9]

So far we have only considered group outcomes. Individual quantities, however, are also quite close to the Cournot Nash prediction of $q_i^N = 19.8$. Figure 2 shows the frequencies of individual quantity choices over all periods. The modal choice in both treatments is at the bracket containing the Nash outcome.

[place Figure 1 about here]

We can analyze individual learning behavior further using the following hit ratios. Let

$$z_i^t := \frac{q_i^t - q_i^{t-1}}{a_i^{t-1} - q_i^{t-1}},$$

where a^{t-1} is the point prediction implied by playing myopic best reply (BR), fictitious play (FP), or imitate the average (AV), respectively. Obviously, $z_i^t = 1$ follows in case of perfect adjustment, while $z_i^t < 0$ implies a severe qualitative violation.[10] Table 3 shows the relative frequency distribution of the z–values.

[9]Thus, there is a distinct difference between our results and those of Cox and Walker (1998) as they find a significant difference between a treatment that should theoretically converge, if best replies are used, and a treatment which should not converge. Further experiments are required to sort out this difference.

[10]In case of $a_i^{t-1} - q_i^{t-1} = 0$ and $q_i^t - q_i^{t-1} \neq 0$ we set $z_i^t < 0$. If also $q_i^t - q_i^{t-1} = 0$, then we set $z_i^t = 1$.

Some observations are immediate from Table 3. While in the majority of cases subjects do adjust in the direction of BR (and to some lesser extent, in the direction of AV), complete adjustments are rare. Only in about 5% of cases do subjects adjust completely to the best reply. Fictitious play, on the other hand, hardly seems to capture behavior adequately since for both treatments almost 50% of adjustments are in the wrong direction. Overshooting ($z > 1.2$) is also more frequent than for the other adjustment rules.

These findings might be compatible with a model in which subjects adjust partially in the direction of BR (see also Rassenti et al., 2000),

$$q_i^t = \gamma_0 + \gamma_1 q_i^{t-1} + \gamma_2 r_i^{t-1}, \tag{17}$$

where r_i^{t-1} denotes subject i's best reply given the other firms' quantities in $t-1$ and $\gamma_1, \gamma_2 > 0$. However, this model cannot explain the 21.5 or 26.4 % of cases in which adjustment pointed away from BR ($z < 0$).

Therefore, we have estimated a model which allows for a combination of adjustments to BR, FP, and AV,

$$q_i^t - q_i^{t-1} = \beta_0 + \beta_1 \left(r_i^{t-1} - q_i^{t-1} \right) + \beta_2 \left(i_i^{t-1} - q_i^{t-1} \right) + \beta_3 \left(f_i^{t-1} - q_i^{t-1} \right), \tag{18}$$

where f_i^{t-1} is the best reply against fictitious play beliefs (given in (8)), and i_i^{t-1} denotes the average quantity of the other firms' output in $t-1$. Note that a subject who strictly played a myopic best reply every period would have $\beta_1 = 1$ and $\beta_k = 0, k \neq 1$. Analogously, $\beta_3 = 1$ for someone who follows fictitious play and $\beta_2 = 1$ for someone who always imitates the average.

We have estimated (18) with weighted least squares (WLS) using $t^{0.2}$ as a weight for all observations.[11] To account for individual differences in learning behavior we added subject intercept and slope dummies.[12] Table 4 shows the results of the regression with pooled data of all subjects.

The results for both treatments are quite similar. The coefficients β_1 and β_2 are significant at the 1% level in both treatments. The coefficients for fictitious play are not significantly different from zero in treatment A, and significant only at the 10% level in treatment B. Given the size of the coefficients it seems that subjects played a mixture of best reply and "imitate

[11] Goldfeld–Quandt tests indicated that variances were significantly lower in later rounds, which can be expected if learning behavior converges. To correct for heteroscedasticity we used WLS with weights chosen so as to maximize the log–likelihood function.

[12] Using a "backward" procedure to eliminate insignificant dummies.

Table 4: WLS Regressions with pooled data

Treatm.	β_1	β_2	β_3	β_0	R^2	DW	Obs.
A	.428** (.059)	.315** (.045)	.037 (.089)	1.19** (.231)	.49	1.97	936
B	.264** (.063)	.148** (.044)	.173* (.090)	.293 (.458)	.46	1.99	626

Note: * significant at 10%, ** significant at 1% level; standard deviations in parentheses. DW = DurbinWatson statistic. Only periods in which subjects were allowed to adjust their quantities are included.

the average", which is evidence in favor of the 'mixed process' analyzed in Section 3.3.

Doing simulations, we found interesting qualitative properties of the mixed process (13) and fictitious play. Both converge very quickly to the Cournot equilibrium. Take, as an example, the values of β_1 and β_2 in treatment A, but normalized such that they add up to one: $\alpha = 0.428/(0.428 + 0.315) \approx 0.57$. For this value, the mixed process would converge to a 1% interval of the Cournot equilibrium values in only 6.5 periods on average.[13] Similarly, fictitious play takes between 4 and 12 periods to converge to this 1% interval. On the other hand, the pure best reply process (with inertia) converges in 25.8 periods to the 1% interval on average. For a typical simulation of these processes see Figure 2. While a pure best reply process explodes, the best reply process with inertia converges to the Nash equilibrium but convergence is slow. The fictitious play as well as the mix of best reply and imitation converge much faster.

Finally, we can compare the theoretical autocorrelations given in Table 1 with the empirical ones shown in Table 2. Clearly, neither a pure best reply process nor a pure imitation process comes even close to the empirical autocorrelations.[14] However, the mixed process does much better if we use for α the estimated values from Table 4. Likewise, fictitious play produces autocorrelations quite close to what is observed in the experiment. For large t, the theoretical autocorrelations with FP converge to 0 in treatment A and to 1/3 for treatment B.

[13] For starting values uniformly distributed between 0 and 100.

[14] A similar observation is made by Rassenti et al. (2000) who conducted oligopoly experiments with five firms. They find convergence at the aggregate level, but autocorrelation is positive, too. Like us they conclude that a pure best-reply process cannot explain the data.

5 Conclusion

In this paper, we report results of an experiment designed to test whether the best reply adjustment process causes unstable markets as predicted by Theocharis (1960). We find no sign of instability. Play converges roughly to the Cournot equilibrium prediction in both treatments, whereas a best-reply process would predict stability in treatment A and explosive fluctuations in treatment B.

We explore several explanations for our result based on alternative learning processes. In particular, we test a basic version of fictitious play and a process based on a mix between playing best replies and imitating the average quantity of the other firms. While fictitious play captures the overall properties of our data quite well (convergence to Nash equilibrium and the autocorrelations), it does not capture the individual round by round decisions of subjects as shown by the very low hit rates and insignificant coefficients in the regressions. The regression results rather support the mixed process between best reply and imitation. Of course, one should not judge such an ad hoc learning theory based on just one experiment. Further experimental research is needed to assess the explanatory power of the mixed process.

References

[1] Abbink, K. and Sadrieh, K. (1995) "RatImage: Research Assistance Toolbox for Computer Aided Human Behavior Experiments", Discussion Paper B-325, University of Bonn.

[2] Asch, S. (1952), Social Psychology, Englewood Cliffs: Prentice Hall.

[3] Cournot, A. (1838) Reserches sur les Principles Mathematiques de la Theorie des Richesses, Paris: Hachette, translated as Research into the Mathematical Principles of the Theory of Wealth, New York: Kelley, 1960.

[4] Cox, J.C. and Walker, M. (1998) "Learning to play Cournot Duopoly Strategies", Journal of Economic Behavior and Organization 36, 141-61

[5] Fisher, F. (1961), "The Stability of the Cournot Solution: The Effects of Speeds Adjustment and Increasing Marginal Costs", Review of Economic Studies, 74,125-135.

[6] Gale, D. and Rosenthal, R. (1999), "Experimentation, Imitation, and Stochastic Stability", Journal of Economic Theory 84, 1-40.

[7] Holt, C. (1985), "An Experimental Test of the Consistent–Conjecture Hypothesis", American Economic Review, 75, 314-325.

[8] Huck, S., Normann, H., and Oechssler, J. (1999), "Learning in Cournot Oligopoly: An Experiment", Economic Journal 109, C80-95.

[9] McManus, R. and Quandt, R.E. (1961), "Comments on the Stability of the Cournot Oligopoly Model", Review of Economic Studies, 74, 136-139.

[10] Monderer, D. and Shapley, L. (1996), "Potential Games", Games and Economic Behavior, 14, 124-143.

[11] Rassenti, S., Reynolds, S., Smith, V., and Szidarovszky, F. (2000), "Adaptation and Convergence of Behavior in Repeated Experimental Cournot Games", Journal of Economic Behavior and Organization 41, 117-146.

[12] Robinson, J. (1951), "An Iterative Method of Solving a Game", Ann. Math. 54, 296-301.

[13] Theocharis, R. (1960), "On the Stability of the Cournot Solution of the Oligopoly Problem", Review of Economic Studies, 73, 133-134.

[14] Vega–Redondo, F. (1997), "The Evolution of Walrasian Behavior", Econometrica, 65, 375-384.

Appendix A: Translation of instructions

Welcome to our experiment. Please read these instructions carefully. In the next 1 or 2 hours you will have to make some decisions at the computer. You can earn some real money. But please be quite during the entire experiment and do not talk to your neighbors. Those who do not follow this rule will have to leave and will not get paid. If you have a question please raise your arm.

You will receive your payment discretely at the end of the experiment. We guarantee anonymity with respect to other participants and we do not record any information connecting your name with your performance.

You can operate the computer with the keyboard or the mouse. Before the experiment there is enough time to make yourself familiar with the computer in a trial round. Money in the experiment is denominated in "Taler". At the end we exchange your earnings into DM at a rate of 500 T = 1 DM. The experiment is divided into several rounds. As said we start with a trial round. The real experiment starts with round 1.

You represent a firm which produces and sells a certain product. Besides you there are 3 other firms which produce and sell the same product. Your task is to decide how much to produce of your good. The capacity of your factory allows you to produce between 0 and 100 units each round. Production cost are 1T per unit. All units (also those of the other firms) are sold on a market (like on a stock exchange or in an auction).

For this the following important rule holds: The price can be between 100T and 0T. The more is sold on the market in total, the lower is the price one obtains per unit. To be precise the price falls by 1T for each additional unit supplied. If – this is only an example – the other firms supply together 10 units and your firm supplies 3 units, then total quantity is 13. The resulting price is $100 - 13 = 87$. If the total quantity were 90, the price would be $100 - 90 = 10$. Profit per unit is the difference between the price and the cost per unit of 1T. Note that you make a loss if the price is lower than the per unit cost. Your profit in a given round results from multiplying the profit per unit with your supplied quantity.

In each round the quantities of all firms are recorded and the resulting profits are calculated. In each round you will be told your profit. Profits from all periods are added and the sum is paid out to you in cash at the end. Additionally you receive a fixed payment of 150T each round. This will be added to your profit each round.

In the first round you decide on a quantity you want to produce and sell. In all further rounds chance decides whether you have the opportunity to revise your quantity. The computer has a mechanism which is comparable to a "one–armed bandit": If you draw a "1" or a "2", you may change your quantity. If you draw a "0", you may not. That is, you may change your quantity in 2 out of 3 cases. With a "0" the quantity of last period is supplied automatically again. Note, that your quantity might be fixed for

This ¶ for treatment B only.

several rounds. Following a “1” or a “2” you may revise your quantity.

In this case you will receive the following information. You are told the total quantity of the other firms last period, and last period’s price.

Additionally, you have access to a profit calculator. The profit calculator is shown on the last page of the instructions. It has two functions: 1. It calculates your profit for arbitrary quantity combinations. That is, you can enter two values, a total quantity for the others (button “A”) and a quantity for yourself (button “I”), and the machine tells you how much you would earn. 2. You can let it calculate for arbitrary quantities of others (button “A”) the quantity at which you would make the highest profit (button “M”). You can use the machine as much as you want before each decision. Before we start you will have enough time to get to know the profit calculator directly at the computer.

Everything we have explained to you holds for the other firms as well. In fact, you are all reading exactly identical instructions.

The experiment lasts for 40 periods in total. Afterwards you will receive your payments in DM. We want to reassure you again that all data will be treated confidentially.

Appendix B: screenshot

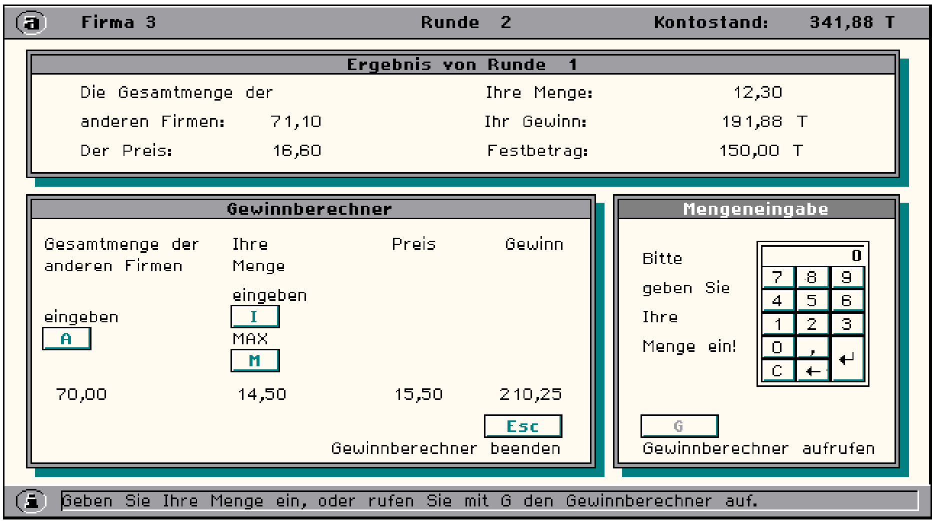

Figure 3: Screenshot

Translation (from top to button, left to right):

Bar at top: Firm 3, Round 2, Balance: 341.88 T

Window at top: Result of round 1, Total quantity of other firms: 71.10, The price: 16.60, Your quantity: 12.30, Your profit: 191.88 T, Fixed payment: 150.00 T.

Lower left window: Profit calculator, Enter total quantity of other firms, Enter your quantity, Price, Profit, Exit profit calculator: Esc.

Lower right window: Enter quantity, Please enter your quantity, open profit calculator.

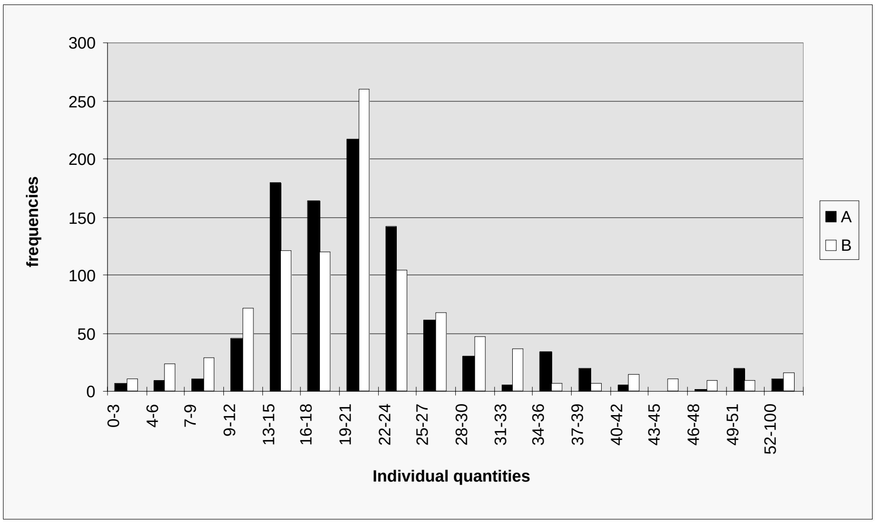

Figure 1: Frequencies of individual quantities: treatment A and B

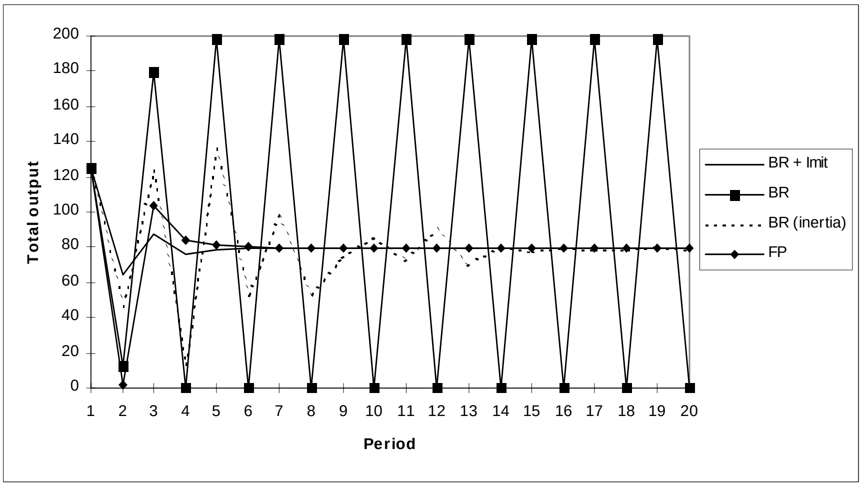

Figure 2: BR + Imit shows the simulation of equation (11) with $\alpha = .57$. BR is the regular best reply process, and BR with inertia is the best reply process with $\theta = \frac{1}{3}$. FP is ficticious play.