Stability of the Cournot Process - Experimental Evidence

Steffen Huck
Department of Economics
Humboldt University, Berlin

Hans–Theo Normann
Department of Economics
Humboldt University, Berlin

Jörg Oechssler*
Department of Economics
Humboldt University, Berlin

July 22, 1997

Abstract

We report results of a series of experiments designed to test the stability of the best reply process. With linear demand and cost functions, the process is stable if and only if there are less than three firms in the market. However, we find no experimental evidence of such instability in a four firm oligopoly. Moreover, there are no differences between a market which theoretically should not converge to Nash equilibrium and one which should converge because of inertia.

JEL– classification numbers: C72, C92, L13

Very preliminary – comments welcome

*Institut für Wirtschaftstheorie, Spandauer Str. 1, 10178 Berlin, Germany, FAX: +49–30–2093-5619, E–mail: oechsler@wiwi.hu-berlin.de

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 Cournot solution does not converge: if there are three firms, finite oscillations about 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 very 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 instantaneous and perfect adjustment is possible. In treatment B, firms must stick to last period's quantity with a probability of 1/3. In Huck, Normann and Oechssler (1997) we show that such a system with inertia is stable. Thus, the theoretical 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 our two treatments. In both cases play converged to the Cournot prediction. Average quantities were slightly above, but still rather close to those of the Cournot equilibrium. In experimental oligopoly markets, Theocharis' (1960) instability result does not occur.

Our results are in contrast to those in Cox and Walker (1997). They also analyse 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 neighbourhood of the equilibrium (with firm 1's quantity on the horizontal axis). Cox and Walker (1997) show that there is a sharp distinction between those two cases supporting the theoretical results. Play was almost never near the theoretically unstable equilibrium but converged nicely to the theoretically stable equilibrium.

To our knowledge, there have been only few experimental studies of Cournot competition with three or more firms. We think these studies do not give definite evidence whether or not Cournot's best reply process is stable because none of them compares the results of a treatment which is theoretically unstable to one which is. Fouraker and Siegel (1963) ran some duopoly and tripoly experiments. They notice a greater variance of outputs in the tripoly markets as compared to duopoly. Fouraker and Siegel (1963, p.265) attribute this to the instability of the best reply process. It is, as we will show, very difficult to draw this conclusion from the greater variance of the tripoly sessions alone. Rassenti *et al.* (1996) ran oligopoly experiments with five firms. They observed no convergence at all. However, Rassenti *et al.* (1996, p.22) explicitly claim that this cannot be attributed to the instability in the sense of Theocharis (1960).

2 Experimental design

In a series of computerised[1] 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

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

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 number of periods was 40 in all sessions and this was commonly known. Subjects possessed all essential information about the market, i.e.

[1]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.

they were informed about the symmetric demand and cost functions in plain words.[2] 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 she could either enter some amount as her own quantity in which case the calculator informed her about the resulting price and her resulting personal profit. Or, she could press a "Max"–button in which case she was informed about 'the quantity which would yield her the highest payoff given the total amount of others'. Additionally, the calculator computed price and profit for this best response.[3] 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 its quantity decision. This was done by a "one–armed bandit" which appeared on 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. Among the economics and business students most did not have any prior knowledge in oligopoly theory.

In each session eight subjects participated. 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

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

[3]In the experiments we did not use the phrase 'best response.'

the experiments.

Subjects were paid according 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 (1960) 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.[4] So additionally subjects earned a fixed payoff of Taler 150 each round. The average payoff was about DM 37.84 which is 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) were introduced and could be tested. When subjects were familiar with both, the rules and the handling of the computer programme, we started the first round.

3 Theoretical Predictions

We start by reproducing Theocharis' (1960) result for the parameters of our model. The behaviour of the four firms is, in matrix notation, given by the following system of equations:

$$\mathbf{q}^t = \mathbf{a} + \mathbf{b}\mathbf{q}^{t-1}, \tag{5}$$

where $\mathbf{q}^t$ is the vector of outputs in t, $\mathbf{a}$ is a four component column vector, whose elements are $99/2$, and $\mathbf{b}$ is a 4×4 matrix with zeros on the diagonal and $-1/2$ off the diagonal. Since $\mathbf{a}$ is a constant and $\mathbf{b}$ is non-singular, the stability of this difference equation depends only on the eigenvalues of $\mathbf{b}$. These are $1/2$ with multiplicity 3 and $-3/2$. The difference equation is not stable, because the latter eigenvalue exceeds 1 in absolute value. The system explodes: players start exceeding the Cournot equilibrium quantities and eventually oscillate between zero output and the monopoly output.

In Huck, Normann and Oechssler (1997), we show that the process with inertia converges globally to the static Cournot equilibrium. The proof is based on the theory of potential games (Monderer and Shapley, 1996)

[4] 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.

and proceeds by proving the existence of an improvement path from any arbitrary state to the equilibrium.

We consider the autocorrelation in quantities an important indicator whether and to what extent the above instability prediction holds. If θ denotes the probability that a firm must stick its quantity, we have

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

and so

$$Q^t = 2(1-\theta)99 - (1-\theta)Q^{t-1}. \quad (7)$$

Thus, there should be negative autocorrelation as long as $\theta \leq \frac{1}{2}$. In treatment A we have $\theta = 0$ and in treatment B we have $\theta = \frac{1}{3}$. Hence, theory predicts a negative autocorrelation for both treatments.

As shown in Huck, Normann and Oechssler (1997) *imitation* plays an important role in oligopoly. As subjects could not observe individual quantities, the only possible form of imitation is to imitate the average quantity of the three other firms, which would result in the following process:

$$q_i^t = \frac{Q_{-i}^{t-1}}{3}. \quad (8)$$

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 the solution to the above difference equation shows

$$q_i^t = \frac{3q_i^1 - Q_{-i}^1}{4} \cdot (-\frac{1}{3})^t + \frac{Q^1}{4}. \quad (9)$$

With inertia the process depends on the realisations of the randomisation device and is therefore path dependent.

4 Experimental results

Figures 2 and 3 in Appendix B show the development of aggregate quantities in treatments A and B over time. The horizontal solid line represents the aggregate Cournot equilibrium output $4q^N = 79.2$. A casual look at these figures reveals that there does not seem to be a noticeable difference between the two treatments. In both treatments play converges fairly well to the

Table 1: Summary Statistics

Treatment	Mean 35	Mean 20	SDev 35	SDev 20	Autocorr.
A	82.83 (6.17)	83.98 (6.79)	8.92 (4.07)	8.15 (4.85)	−.043
B	84.32 (4.56)	82.56 (2.48)	13.11 (11.94)	10.00 (10.64)	.387

Note: Mean35 (Mean20) is the average total quantity and SDev35 (SDev20) is the average standard deviation of total quantities, both measured over the last 35 (20) periods and over all six groups in one treatment. Standard deviations in parentheses.

Table 2: OLS Regressions with pooled data

Treatment	β_1	β_2	β_0	R^2	DW	Obs.
A	.437 (.028)	.324 (.030)	.972 (.215)	.40	1.93	936
B	.430 (.038)	.340 (.038)	1.42 (.377)	.41	2.05	610

Note: Standard deviations in parentheses. DW = DurbinWatson statistic. Subject dummies are used with the restriction that their coefficients sum to zero. Only periods in which subjects were allowed to adjust their quantities are included.

Cournot output. Note that there was no successful attempt to collude: the aggregate output was only twice at, or below, the collusive output of 49.5.

Table 1 reports summary statistics for the last 35 and the last 20 periods, respectively. 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 differences between the mean quantities for treatment A and B at any reasonable significance level. Intuitively, one should think that at least the variability of outputs is larger in treatment A, since the inertia should have a damping effect. But, surprisingly, standard deviations are even higher in treatment B. This is in contrast to Fouraker and Siegel's (1963, p. 265) observation. The average autocorrelation coefficients of Q^t are slightly negative for treatment A. For treatment B the inertia seems to induce a positive autocorrelation. Both coefficients are larger than the theory predicts if best replies were being played.

We have estimated with OLS the following equation

$$q_i^t - q_i^{t-1} = \beta_0 + \beta_1 \left(r_i^{t-1} - q_i^{t-1} \right) + \beta_2 \left(i_i^{t-1} - q_i^{t-1} \right) \tag{10}$$

where r_i^{t-1} denotes subject i's best reply (i.e. reaction function) given the other firms' quantities in $t-1$; i_i^{t-1} denotes the average quantity of the other firms' output in $t-1$. Note, that a subject who played strictly a myopic best reply every period would have $\beta_1 = 1$ and $\beta_k = 0, k \neq 1$.[5] Similarly, for someone who follows the rule "imitate the average". Table 2 shows the results of this regression. The results for both treatments are virtually identical.

It seems that firms play a mixture of best reply and "imitate the average". Generally, if α denotes the weight given to best replies and $1-\alpha$ the weight given to the average quantity of the other firms, we get the following difference equation

$$q_i^t = \alpha \frac{99 - Q_{-i}^{t-1}}{2} + (1-\alpha) \frac{Q_{-i}^{t-1}}{3}. \tag{11}$$

This difference equation is stable if $\alpha < \frac{4}{5}$. If it is stable, it converges to the Cournot equilibrium. Doing simulations, we found another qualitative property of this process: it converges extremely quickly. Take, as an example, the values of β_1 and β_2 in treatment A, but normalised such that they add up to one: $\alpha = 0.437/(0.437 + 0.324) \approx 0.57$.For this value, play would converge to a 1% interval of the Cournot equilibrium values in only 6.5 periods on average. The pure best reply process (with inertia),converges in 25.8 periods to 1% interval on average. For an illustration of the simulations of these processes, see Figure 4.

5 Conclusion

In this paper, we report results of a series of experiments designed to test the stability of the best reply adjustment process. We find no sign of instability, play converges roughly to the Cournot equilibrium prediction. Somewhat surprisingly, the experimental market which is theoretically unstable is less variable than the market which should converge in theory. We present a

[5] Thus, one advantage of the difference to q_i^{t-1}–form of the regression is that the coefficients have a nice interpretation. The other advantage is that it avoids problems of serial correlation.

possible explanation for our result by suggesting that subjects do not play a pure best reply but mix between playing best replies and imitating the average quantity of the other firms.

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] 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.

[3] COX, J.C. and WALKER, M. (1997) "Learning to play Cournot Duopoly Strategies" *mimeo (University of Arizona).*

[4] 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.

[5] FOURAKER, L. and SIEGEL, S. (1963) *Bargaining Behavior*, New York: McGraw–Hill.

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

[7] HUCK, S., NORMANN, H., and OECHSSLER, J. (1997), "Stability of the Cournot Process – Experimental Evidence", mimeo (Humboldt University, Berlin).

[8] MCMANUS, R. and QUANDT, R.E. (1961), "Comments on the Stability of the Cournot Oligopoly Model", *Review of Economic Studies*, 74, 136-139.

[9] MONDERER, D. and SHAPLEY, L. (1996), "Potential Games", *Games and Economic Behavior*, 14, 124-143.

[10] RASSENTI, S., REYNOLDS, S., SMITH, V., and SZIDAROVSZKY, F. (1996), "Adaptation and Convergence of Behavior in Repeated Experimental Cournot Games", *mimeo (University of Arizona).*

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

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

This ¶ for treatment B only.

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.

Have fun!

Appendix B: screenshot and figures

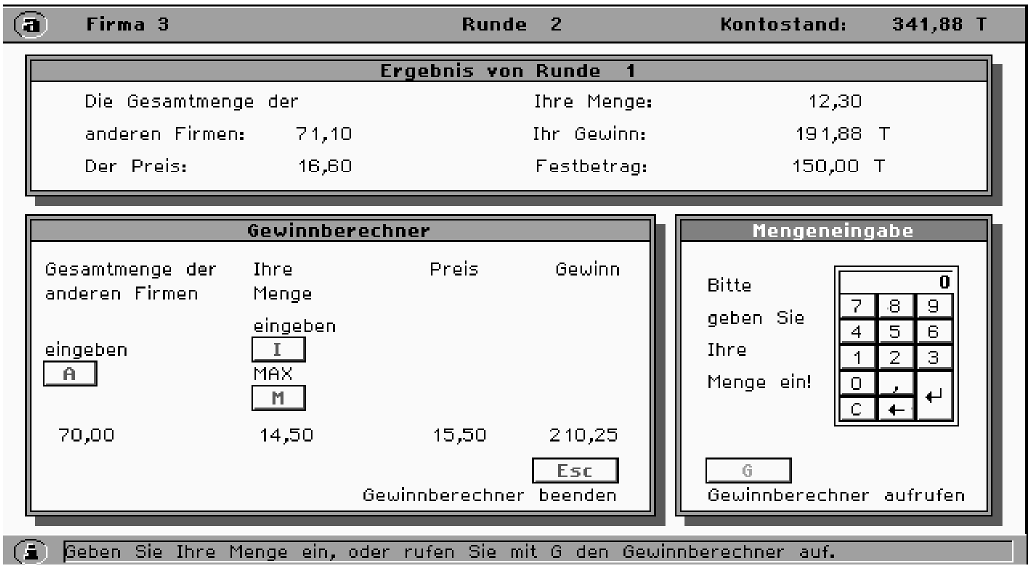

Figure 1: Screenshot

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.

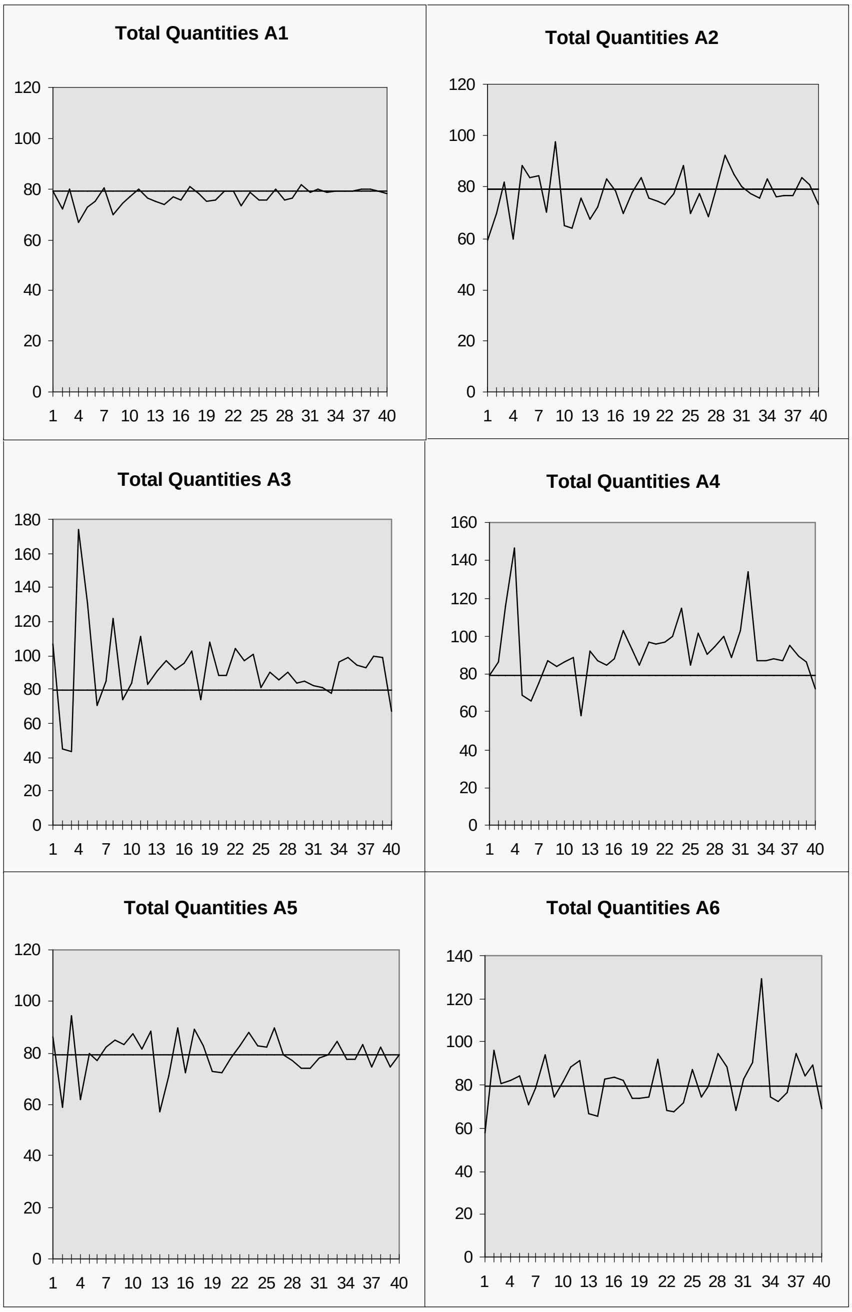

Figure 2: Treatment A. Solid line is the Nash quantity $Q^N = 79.2$.

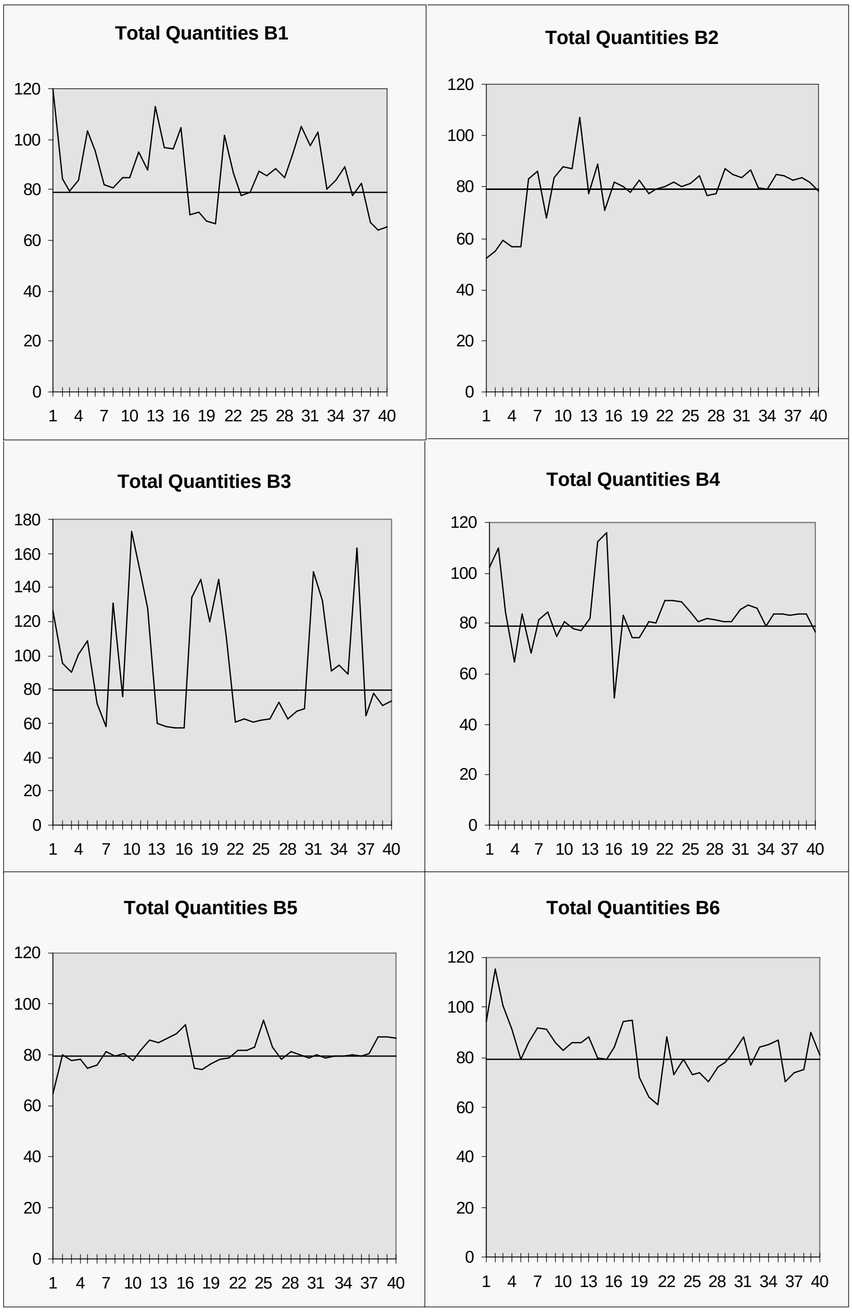

Figure 3: Treatment B. Solid line is Nash quantity $Q^N = 79.2$.

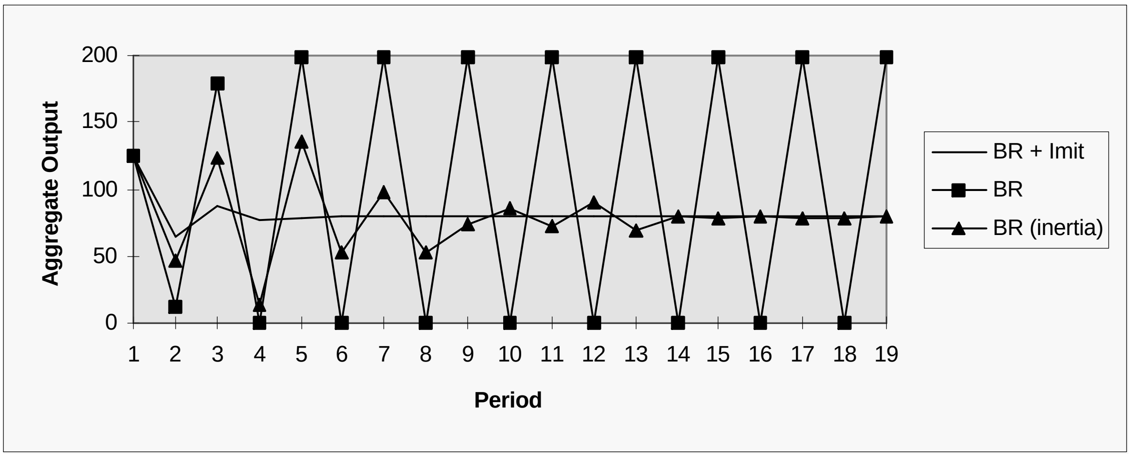

Figure 4: BR + Imit shows the simulation of equation (11) with $\alpha = .56$. BR is the regular best reply process, and BR with inertia is the best reply process with $\theta = \frac{1}{3}$.