

Huck, Steffen; Normann, Hans-Theo; Oechssler, Jörg

Working Paper

Through Trial & Error to Collusion - The Discrete Case

Bonn Econ Discussion Papers, No. 6/2000

Provided in Cooperation with:

Bonn Graduate School of Economics (BGSE), University of Bonn

Suggested Citation: Huck, Steffen; Normann, Hans-Theo; Oechssler, Jörg (2000) : Through Trial & Error to Collusion - The Discrete Case, Bonn Econ Discussion Papers, No. 6/2000, University of Bonn, Bonn Graduate School of Economics (BGSE), Bonn

This Version is available at:

<https://hdl.handle.net/10419/78373>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

BONN ECON DISCUSSION PAPERS

Discussion Paper 6/2000

Through Trial & Error to Collusion - The Discrete Case

by

**Steffen Huck, Hans-Theo Norman
and Jörg Oechssler**

May 2000



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

Through Trial & Error to Collusion – The Discrete Case*

Steffen Huck
Department of Economics
Royal Holloway College, London

Hans-Theo Normann
Department of Economics
Humboldt University, Berlin

Jörg Oechssler[†]
Department of Economics
University of Bonn

May 11, 2000

Abstract

In this note we study a very simple trial & error learning process in the context of a Cournot oligopoly. Without any knowledge of the payoff functions players increase, respectively decrease, their quantity by one unit as long as this leads to higher profits. We show that despite the absence of any coordination or punishing device this process converges to a collusive outcome.

JEL- classification numbers: C72, L13.

*We thank seminar audiences at the Stockholm School of Economics, Tilburg University, University of Amsterdam (Tinbergen Institute), the Institute for Advanced Studies, Vienna, Texas A&M, Royal Holloway, and Humboldt, and in particular Tilman Börgers, Ulrich Kamecke, Georg Nöldeke, Roy Radner, Klaus Ritzberger, and Jörgen Weibull for helpful comments. The first author acknowledges financial support from DFG grant Hu 787/1-1.

[†]Adenauerallee 24, 53113 Bonn, Germany, E-mail: oechssler@uni-bonn.de

1 Introduction

In this note we consider a very simple learning process, which we call trial & error learning. The process is simple in two ways. First, it requires a very low cognitive effort of players. And second, it does not require any information about the payoff function of the game. It works as follows. Players choose their strategies from a finite one-dimensional grid. Typical examples include prices, quantities, expenditures and so on. Everytime players change their strategy by one grid point, they check whether this increases or a decreases payoffs. If it increases payoffs, the movement in this direction is continued. If it does not, the reverse direction is taken.

We study the consequences of this learning process in the context of a standard Cournot oligopoly. Somewhat surprisingly, it turns out that trial & error learning yields a collusive outcome.¹ We prove this result analytically for the duopoly case. With simulations we demonstrate that the result is robust to the introduction of more firms and slight modifications of the learning rule.

This result seems remarkable as players cannot observe the actions of their opponents. Thus, there is no basis for any coordination or punishing device. The result is also remarkable as practically all learning processes in the literature converge to the (unique) Cournot–Nash equilibrium of the game if they converge. This holds for best reply learning (Cournot, 1838), fictitious play, evolutionary dynamics like the replicator dynamics, gradient learning (Arrow and Hurwicz, 1960), or more generally for the class of adaptive learning processes (Milgrom and Roberts, 1991). One interesting exception is a simple imitation process, which has been shown to converge to the competitive outcome (Vega–Redondo, 1997).

We do not know of any previous analysis of such trial & error learning processes. Probably most closely related is gradient learning (Arrow and Hurwicz, 1960), which requires, however, much more information. Trial & error learning could be considered a particular form of learning direction

¹By collusive outcome we mean an outcome on the Pareto frontier.

theory (Selten and Buchta, 1998). This theory assumes that players have a model which allows them to conclude in which direction better actions can be found. In the absence of information about demand and cost conditions, one interpretation is that the right direction can be found by determining which direction was successful last period.

The remainder of the paper is organized as follows. The next section introduces the learning rule and presents the main theoretical result. Section 3 contains the simulation results. In the concluding Section 4 we discuss some experimental evidence.

2 Trial & error learning

Consider a standard Cournot oligopoly with n firms. Each firm may choose outputs from a finite grid

$$\Gamma := \{0, \delta, 2\delta, \dots, v\delta\}$$

for some arbitrarily small grid size $\delta > 0$ and $v \in \mathbb{N}$ large enough. Inverse demand, $p(Q) \geq 0$, $Q = \sum_{i=1}^n q_i$ is assumed to satisfy $p' < 0$ and $p' + 2p''Q < 0$. Note that the latter assumption is weaker than requiring concavity of demand.

Firm i 's cost function $C_i(q_i)$ is increasing and weakly convex, that is, $C_i' > 0$ and $C_i'' \geq 0$. Furthermore, we assume that for all i there exists a finite Q such that $p(Q) = C_i'(0)$. In particular, this is satisfied if there is a finite Q such that the price becomes zero. To avoid a monopolized market we assume that each firm's monopoly price, p_i^m , is larger than the minimal marginal cost of each firm, $p_i^m > C_j'(0)$, for all i and j . Let $\Pi_i(q_1, \dots, q_n) = p(Q)q_i - C_i(q_i)$ denote profit of firm i .

Next, let us define the set of collusive outcomes (i.e. the Pareto frontier)

$$\left(\begin{array}{c} \mathfrak{X} \\ \mathbf{q}^c \in \mathbb{R}_+^n : \mathbf{q}^c = \arg \max_{\mathbf{q}} \sum_{i=1}^n \lambda_i (p(Q)q_i - C_i(q_i)), \lambda_i \geq 0, \text{ some } \lambda_j > 0 \end{array} \right) .$$

The joint profit maximum is found by setting $\lambda_i = 1$, for all i .

We assume that players behave according to the following trial & error learning process

$$q_i^t = \max^{\text{©}} \left[0, q_i^{t-1} + \delta s_i^{t-1} \right]^{\text{a}}, \quad (1)$$

where

$$s_i^t := \text{sign}(q_i^t - q_i^{t-1}) \times \text{sign}(\pi_i^t - \pi_i^{t-1})$$

if $(q_i^t - q_i^{t-1}) \cdot (\pi_i^t - \pi_i^{t-1}) \neq 0$.² Otherwise $s_i^t = +1, 0$, or -1 , each with positive probability. In period $t = 0$ players start with some arbitrary q_i^0 and s_i^0 .

This formalizes the idea that a direction (one grid point up or down) that has proved successful last period is repeated again this period.³ If the change in profits or the change in quantities was zero, the quantity either remains the same, is increased by one unit or is decreased by one unit, each with positive probability.

Note that neither information about other firms' demand or cost function, nor information about their past actions, nor information about one's own demand or cost functions are required for the trial & error process to work.

Each plausible learning process should be subject to some noise as individuals generally make mistakes in the execution of their strategies or (more or less) systematically try out different actions. We assume therefore that with some small probability $\varepsilon > 0$ each firm chooses an arbitrary direction of change s_i^t (each possible direction is chosen with some fixed, positive probability). At the cost of some more cumbersome notation we could equally assume that this probability is different for each firm and/or time period or that deviations to quantities which are further away than one grid point are possible.

Our assumptions define a Markov process on the finite state space $\Gamma^n \times \{-1, 0, 1\}^n$, where a state is given by $(\mathbf{q}^t, \mathbf{s}^t)$. For $\varepsilon > 0$ the process is irreducible and, therefore, has a unique stationary distribution. Formally,

²Note that the upper bound $v\delta$ can never be reached if v large enough.

³The restriction to movements of one grid point may be justified by convex adjustment costs, which can make it optimal to adjust in very small steps.

we consider the limit distribution for $\varepsilon \rightarrow 0$. For $\varepsilon = 0$ the process may have several absorbing sets.⁴ By standard arguments (see e.g. Samuelson, 1994) only the members of absorbing sets of the unperturbed process can appear in the support of the limit distribution. In this paper we will speak of convergence to some point \mathbf{q} if all states \mathbf{q}' , which have positive probability under the limit distribution, are close to \mathbf{q} in the sense that $\|\mathbf{q} - \mathbf{q}'\|$ is of the order of the grid size δ , i.e. $\|\mathbf{q} - \mathbf{q}'\| = O(\delta)$. In other words, by letting the grid size δ go to zero the distance from any point in a limit cycle to \mathbf{q} can be made arbitrarily small.

Theorem 1 For a duopoly the trial & error process converges to a collusive outcome \mathbf{q}^c . If cost functions are identical, $C_1(\cdot) = C_2(\cdot)$, then it converges to the joint profit maximizing outcome.

Proof. We use the notation $\downarrow\uparrow$ to indicate that firm 1 decreases its quantity and firm 2 increases its quantity. Analogously, for $\uparrow\downarrow, \uparrow\uparrow$ and $\downarrow\downarrow$. Let us first consider the change in individual profits due to $\uparrow\uparrow$.

$$\Delta\Pi_i(\uparrow\uparrow) = (q_i + \delta)p(q_1 + q_2 + 2\delta) - q_i p(q_1 + q_2) - C_i(q_i + \delta) + C_i(q_i).$$

We can now implicitly define two functions $z_i(q_i)$, $i = 1, 2$ by

$$\frac{1}{\delta}\Delta\Pi_i(\uparrow\uparrow)(q_i, z_i(q_i)) \equiv 0. \quad (2)$$

We call $z_i(q_i)$ the improvement frontier for firm i since for all $q_j < z_i(q_i)$ ($q_j > z_i(q_i)$) profits of firm i increase (decrease) when both firms raise their quantities. Similarly, we can derive an improvement frontier for $\Delta\Pi_i(\downarrow\downarrow) = 0$, which is one grid point above and to the right of the curve for $\Delta\Pi_i(\uparrow\uparrow) = 0$.

Note that for $\delta \rightarrow 0$ equation (2) becomes

$$\frac{\partial\Pi_i(q_i, z_i(q_i))}{\partial q_i} + \frac{\partial\Pi_i(q_i, z_i(q_i))}{\partial q_j} = 0. \quad (3)$$

⁴A set of states is called absorbing if there is zero probability to exit the set and a positive probability of moving from any state in the set to any other state in the set in finite time.

or equivalently,

$$p(q_i + z_i(q_i)) + 2q_i p'(q_i + z_i(q_i)) - C_i'(q_i) = 0. \quad (4)$$

By implicitly differentiating (4) we find that

$$\frac{dz_i(q_i)}{dq_i} < -1. \quad (5)$$

Next, we define curves r_1 and r_2 in the (q_1, q_2) space along which price equals marginal costs for firm i , i.e.

$$p(q_i + r_i(q_i)) - C_i'(q_i) \equiv 0. \quad (6)$$

Some simple facts are immediate. (i) Implicitly differentiating (6) yields

$$\frac{dr_i(q_i)}{dq_i} \leq -1.$$

(ii) $r_i(0)$ is finite because we have assumed that there exists a finite Q such that $p(Q) = C_i'(0)$. Comparing (4) and (6) shows that (iii) $z_i(0) = r_i(0)$ and (iv) $r_i(q_i) > z_i(q_i)$ for $q_i > 0$.

Next, we show that there exists a unique intersection P_1 of the two improvement frontiers z_1 and z_2 in the interior. Our assumption that $p_i^m > C_j'(0)$, for all i and j , guarantees that $q_i^m < r_j(0) = z_j(0)$. Let q_i^* be defined by $z_i(q_i^*) = 0$. Since q_i^m is defined by $\partial \Pi_i / \partial q_i = 0$, whereas q_i^* is given by (3), we have by the concavity of the profit function that $q_i^* < q_i^m$. Thus, $q_i^* < z_j(0)$. This together with (5) implies that P_1 exists in the interior and is unique (as shown in Figure 1).

[insert Figure 1 about here]

We claim that P_1 corresponds to some q^c . To see this, note that collusion requires for $i = 1, 2; i \neq j$ and some $\lambda_j \geq 0$

$$\frac{\partial \Pi_i}{\partial q_i} + \lambda_j \frac{\partial \Pi_j}{\partial q_i} = 0. \quad (7)$$

Table 1: Movements in $t + 1$

Region	Movement in t			
	$\uparrow\uparrow$	$\uparrow\downarrow$	$\downarrow\uparrow$	$\downarrow\downarrow$
\mathcal{A}	$\uparrow\uparrow$	$\uparrow\uparrow$	$\uparrow\uparrow$	$\uparrow\uparrow$
\mathcal{B}	$\downarrow\downarrow$	$\uparrow\uparrow$	$\uparrow\uparrow$	$\downarrow\downarrow$
\mathcal{C}	$\downarrow\uparrow$	$\uparrow\uparrow$	$\uparrow\uparrow$	$\downarrow\uparrow$
\mathcal{D}	$\uparrow\downarrow$	$\uparrow\uparrow$	$\uparrow\uparrow$	$\uparrow\downarrow$
\mathcal{L}_0	$\downarrow\downarrow$	$\downarrow\downarrow$	$\downarrow\downarrow$	$\downarrow\downarrow$
\mathcal{L}_1	$\uparrow\downarrow$	$\uparrow\downarrow$	$\uparrow\downarrow$	$\uparrow\downarrow$
\mathcal{L}_2	$\downarrow\downarrow$	$\uparrow\downarrow$	$\uparrow\downarrow$	$\downarrow\downarrow$
\mathcal{L}_3	$\downarrow\downarrow$	$\downarrow\uparrow$	$\downarrow\uparrow$	$\downarrow\downarrow$
\mathcal{L}_4	$\downarrow\uparrow$	$\downarrow\uparrow$	$\downarrow\uparrow$	$\downarrow\uparrow$

Thus, conditions (3) and (7) are equivalent if

$$\frac{\partial \Pi_i}{\partial q_j} = \lambda_j \frac{\partial \Pi_j}{\partial q_i}, \quad (8)$$

that is, if $q_i p'(Q) = \lambda_j q_j p'(Q)$ or if $q_i = \lambda_j q_j$. Hence, there always exists some λ_j such that both conditions are equivalent. Clearly, if costs are symmetric, the solutions to (3) and (7) are symmetric, and hence (8) holds for $\lambda_j = 1$. Thus, for $\delta \rightarrow 0$ the improvement frontiers intersect at the joint profit maximum.

Taken together those facts define the shapes of the four curves, r_1, r_2, z_1, z_2 . Figure 1 shows a typical constellation. Note that P_2, P_3 , and P_4 need not necessarily exist.

We can now define nine subsets of the quantity space (some of which might be empty) as shown in Figure 1. In regions \mathcal{L}_0 through \mathcal{L}_4 at least one firm makes a loss.

For each region we can determine the transitions as shown in Table 1. For example, in regions \mathcal{A} through \mathcal{D} $\uparrow\downarrow$ and $\downarrow\uparrow$ are followed by $\uparrow\uparrow$ because $\Delta p = 0$ and $\Pi_i > 0$. In region \mathcal{L}_0 all movements result in $\downarrow\downarrow$ as price is below marginal cost for both firms. In region \mathcal{L}_4 all movements result in $\downarrow\uparrow$ since for firm 1 price is below marginal cost, whereas firm 2 is still below its

improvement frontier. All other entries in Table 1 can be derived similarly. Note in particular that nowhere in regions \mathcal{B} and \mathcal{C} (\mathcal{B} and \mathcal{D}) the movements $\downarrow\uparrow$ ($\uparrow\downarrow$) are possible. Hence, the process in regions \mathcal{B} , \mathcal{C} , and \mathcal{D} weakly moves in the direction of the diagonal.

Applying Table 1 repeatedly it can be seen that after at most two steps in each region directions are given as indicated by the arrows in Figure 1, namely, always decreasing in regions \mathcal{B} and \mathcal{L}_0 , and always increasing in region \mathcal{A} . In regions \mathcal{C} and \mathcal{D} quantities are zigzagging in the direction of region \mathcal{B} . In regions \mathcal{L}_2 and \mathcal{L}_3 both indicated directions are possible depending on the starting condition.

Considering the dynamics in Figure 1 it is clear that the only candidates for absorbing sets are limit cycles around one of the intersections P_1 through P_4 . Especially, it can be checked easily that no cycles of length two between regions \mathcal{B} and \mathcal{C} or between regions \mathcal{B} and \mathcal{D} are possible. Clearly, P_1 is a sink of the dynamics whereas P_2 , P_3 , and P_4 are saddles. The latter are not stable to noise which can always put the process into region \mathcal{B} and thus in the basin of attraction of P_1 . The sink P_1 , however, is robust to noise. Thus, for ε small the process will be almost always in an absorbing set around P_1 . The maximum distance between any point q' in such an absorbing set and q^c is bounded by the fact that after at most two transitions inside each region the directions are as indicated by the arrows in Figure 1. Thus, the distance between q^c and q' is of the order of δ . ■

Remark 1 If the cost functions are linear and the same for both firms, then the intersection points P_2 , P_3 , and P_4 do not exist since in this case $\frac{dr_i(q_i)}{dq_i} = -1$. Therefore, global convergence to the neighborhood of P_1 is assured even without noise.

The intuition for Theorem 1 is the following. It is relatively easy to see why firms which are perfectly aligned will move to a collusive outcome. Suppose two symmetric firms start from some output $q_1 = q_2$ larger than the collusive outcome. If both decrease their quantity, both increase their profits and will continue to do so until the collusive outcome is reached.

This example also shows why the Cournot outcome is not a rest point of our process. Once firms surpass the collusive outcome, profits are lowered by further reductions in quantities. Hence, both firms will turn and start jumping around the collusive outcome. In fact, the typical cycle which emerges in simulations consists of both firms moving in step $(q^c - \delta, q^c - \delta) \rightarrow (q^c, q^c) \rightarrow (q^c + \delta, q^c + \delta) \rightarrow (q^c, q^c) \rightarrow (q^c - \delta, q^c - \delta)$,⁵ though more complicated patterns are possible.

The question then arises why firms which start from arbitrary initial quantities and directions of change (s_i^0) , could become perfectly aligned. Suppose two firms with different quantities move downwards. They will continue to do so until at least one firm's profit decreases. If firms are not too close together, it is always the firm with the smaller output which hits this boundary first. This is because the firm selling the higher quantity gains more from the increase in price. Thus, while the firm with lower output already moves upwards, the other firm continues to move downwards thereby decreasing the distance between the firms by 2 grid points. Similarly, when moving upwards the firm with higher output hits the boundary first. Thus, there is a general tendency to equalize quantities. In terms of Figure 1 firms generally move to the diagonal. Once on the diagonal, the argument of the previous paragraph applies.

Finally, consider a parallel to the conjectural variations model (Hicks, 1935). In this model when the conjectural derivative is +1, the equilibrium conditions are identical to our equation (3). Note, however, that conjectural variations models exogenously fix the conjectural derivative to +1, while with our learning process (3) results endogenously.

3 Simulations results

In order to check whether Theorem 1 is simply an artefact of our assumptions, we ran computer simulations to assess its robustness.⁶ In particular,

⁵This cycle also exists with more than two firms.

⁶The programming was done in Turbo Pascal. The source code is available from the authors upon request.

we simulated oligopolies with up to 10 firms and various functional forms both, for symmetric and for asymmetric firms. Further, we analyzed two modifications of the learning rule.

3.1 n-firm oligopolies

The most important result from the simulations is that the prediction of Theorem 1 also holds for more than two firms. We simulated oligopolies with $n = 2, \dots, 10$ firms and for three functional forms, linear demand and cost, linear demand and quadratic cost, and quadratic demand and linear cost.⁷ In all simulations with symmetric firms, play converged to the joint profit maximizing outcome.⁸ With cost asymmetries a Pareto efficient outcome was reached.

Result In simulations, the trial & error process with noise converged globally to the collusive outcome q^c in all cases. When cost functions were identical, $C_1(\cdot) = C_2(\cdot)$, it converged to the joint profit maximizing outcome.

As pointed out above without noise the process can get stuck in limit cycles around intersection points P which are far away from Pareto efficient outcomes. To see how large the proportion of such limit cycles is for different numbers of firms we also ran simulations without noise. Table 2 shows the average percentage of limit cycles in 10,000 simulations with random starting quantities for oligopolies with linear demand and cost.

As expected, the number of limit cycles increases with the number of firms (and therefore with the number of possible intersection points P). Numerical values for quadratic specifications are similar. No matter how

⁷The functional forms used for symmetric firms were the following: (a) $\Pi_i = (1.1 - q_1 - q_2)q_i - 0.1q_i$; (b) $\Pi_i = (1 - q_1 - q_2)q_i - (q_i)^2$; (c) $\Pi_i = (1.1 - (q_1 - q_2)^2)q_i - 0.1q_i$. For asymmetric firms, different cost parameters were employed. The grid size was in all cases $\delta = 0.001$. Initial quantities were random, but restricted to $q_i < 1$. The noise was modelled such that there is a probability $\varepsilon = 1/500$ for each firm in each period that s_i^t is chosen with equal probabilities from $\{-1, 0, 1\}$.

⁸Here again, convergence to q means convergence to a limit cycle which is close to q .

Table 2: Percentage of limit cycles

Numer of firms	2	3	4	5	6	7	8	9	10
Limit cycles in %	0	2.65	2.69	7.34	7.64	11.59	12.70	17.84	20.16

large the number of limit cycles, with noise play always converged to the collusive outcome.

3.2 Modifications of the learning rule

One property of our trial & error process is that all firms move with the same step size δ . What happens if firms have different step sizes δ_i ? Consider for example a symmetric oligopoly with constant marginal cost in which one firm's step size is twice that of the other firms. What happens in simulations is that total output still converges to the joint profit maximum. However, the distribution of output changes: the firm with the larger step size has twice the output share of the other firms. The intuition is that from the perspective of the remaining firms the first firms moves and behaves like two firms which are perfectly aligned. In fact, we obtained this result for up to 10 firms, and with both, linear and quadratic demand. The factor k by which one firm's step size is multiplied is arbitrary — the first firm will end up with k times the shares of the other firms. Even with quadratic cost there is convergence to some Pareto efficient outcome in which firms produce different quantities (also if costs are asymmetric).

A second feature of our learning process is that only the direction of movement but not the step size is influenced by the change in profits. We ran simulations where instead of using constant step size δ we used an endogenous step size of

$$\delta = \frac{\mu \frac{\Pi_i^{t-1} + c}{\Pi_i^{t-2} + c} \uparrow d}{\cdot}$$

The constant c served to avoid division by zero, and d was chosen to be even in order to avoid changes in sign.⁹ Again we found convergence to a

⁹For c we chose values between 0.1 and 1, and for d between 2 and 8.

collusive outcome.

4 Conclusion

In this note we studied a trial & error learning process which converges to a collusive outcome. This result may be somewhat surprising as players in our setting are totally ignorant of the other players; they do even not realize that they are playing a game. In fact, they do not have to know their own payoff function. Nevertheless, they manage to coordinate on a collusive outcome endogenously.

We believe that the trial & error process has some intuitive appeal. In particular, in situations in which players are not familiar with the payoff structure of the game it seems plausible that they adjust cautiously in a direction that has proved to be successful. Ultimately, it is an empirical question whether players behave according to such a process or not.

There is some experimental evidence, both on the individual level and on the aggregate level. On the individual level Huck, Normann, and Oechssler (1999) show that the direction of change is predicted correctly by the trial & error process for 80% of subjects' choices. However, since the experiment was not designed to test trial & error learning in particular,¹⁰ subjects were not constrained to change their quantity by one grid point only and mostly adjusted by more. Secondly, the theory requires that all subjects play according to the rule. If only some subjects violate it occasionally, one cannot expect convergence to the theoretical prediction any more. Hence, it is not surprising that total quantities were not even close to the collusive outcome. On the aggregate level several other studies found substantial support for a tendency toward collusion but mostly in duopoly cases (see Holt, 1995, for an overview).

An interesting question for further research is to identify a general class of games where trial & error learning yields Pareto efficiency. We have found some examples where this might work. For instance, trial & error learning

¹⁰Rather, the experimental results lead us to analyze the trail & error process.

yields cooperation in a finite mixed extension of the Prisoner's Dilemma, where players move on a probability grid and receive expected payoffs.¹¹

As pointed out in the Introduction the Cournot outcome is a very robust prediction as practically all known learning processes converge to it. Any imaginative reader is sure to find modifications of our trail & error learning rule which do so as well. However, the interesting fact is that there are plausible learning processes which do not converge to the Cournot outcome.

References

- [1] Arrow, K.J., and Hurwicz, L. (1960), "Stability of the Gradient Process in n -Person Games", *Journal of the Society of Industrial and Applied Mathematics*, 8, 280-294.
- [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] Hicks, J.R. (1935), "Annual Survey of Economic Theory: The Theory of Monopoly", *Econometrica*, 3, 1-20.
- [4] Holt, C. (1995), "Industrial Organization: A Survey of Laboratory Results", in: *The Handbook of Experimental Economics* (J. Kagel and A. Roth, eds.), Princeton: Princeton University Press.
- [5] Huck, S., Normann, H., and Oechssler, J. (1999), "Learning in Cournot Oligopoly: An Experiment", *Economic Journal*, 109, C80-C95.
- [6] Milgrom, P. and Roberts, J. (1991), "Adaptive and Sophisticated Learning in Normal Form Games", *Games and Economic Behavior*, 3, 82-100.
- [7] Samuelson, L. (1994), "Stochastic Stability in Games with Alternative Best Replies", *Journal of Economic Theory*, 64, 35-65.

¹¹This example has been suggested by Jörgen Weibull.

- [8] Selten, R. and Buchta, J. (1998), "Experimental Sealed Bid First Price Auctions with Directly Observed Bid Functions", in *Games and Human Behavior: Essays in Honor of Amnon Rapoport* (D. Budescu, I. Erev, and R. Zwick, eds.) Lawrence Erlbaum Ass., Mahwah, NJ.
- [9] Vega-Redondo, F. (1997), "The Evolution of Walrasian Behavior", *Econometrica*, 65, 375-84.

