

Argenton, Cédric; Ivanova-Stenzel, Radosveta; Müller, Wieland

Conference Paper

Cournot Meets Bayes-Nash: A Discontinuity in Behavior in Finitely Repeated Duopoly Games

Beiträge zur Jahrestagung des Vereins für Socialpolitik 2022: Big Data in Economics

Provided in Cooperation with:

Verein für Socialpolitik / German Economic Association

Suggested Citation: Argenton, Cédric; Ivanova-Stenzel, Radosveta; Müller, Wieland (2022) : Cournot Meets Bayes-Nash: A Discontinuity in Behavior in Finitely Repeated Duopoly Games, Beiträge zur Jahrestagung des Vereins für Socialpolitik 2022: Big Data in Economics, ZBW - Leibniz Information Centre for Economics, Kiel, Hamburg

This Version is available at:

<https://hdl.handle.net/10419/264079>

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.

Cournot meets Bayes-Nash: A discontinuity in behavior in finitely repeated duopoly games*

Cédric Argenton[†]

CentER & TILEC, Tilburg University

Radosveta Ivanova-Stenzel[‡]

Technical University Berlin

Wieland Müller[§]

VCEE, University of Vienna,

CentER & TILEC, Tilburg University

14th February 2022

Abstract

We conduct a series of Cournot duopoly market experiments with a high number of repetitions and fixed matching. Our treatments include markets with (a) complete cost symmetry and complete information, (b) slight cost asymmetry and complete information, and (c) varying cost asymmetries and incomplete information. For the case of complete cost symmetry and complete information, our data confirm the well-known result that duopoly players achieve, on average, partial collusion. However, as soon as any level of cost asymmetry or incomplete information is introduced, observed average individual quantities are remarkably close to the static Bayes-Nash equilibrium predictions.

JEL Classification numbers: D43, L13, C72, C92.

Keywords: Cournot, Bayesian game, Bayes-Nash equilibrium, repeated games, collusion, cooperation, experimental economics.

*The experiments reported on in this paper were funded by the University of Vienna. Support by CRC TRR 190 “Rationality and Competition” as well as by the Berlin Centre for Consumer Policies (BCCP) is gratefully acknowledged.

[†]Department of Economics, Tilburg University, Warandelaan 2, Postbus 90153, 5000 LE Tilburg, The Netherlands; E-mail: c.argenton@uvt.nl.

[‡]Department of Economics and Economics Law, School of Economics and Management, Technische Universität Berlin, Str. des 17. Juni 135, 10632 Berlin, Germany; E-mail: ivanova-stenzel@tu-berlin.de.

[§]Department of Economics, VCEE, University of Vienna, Oskar-Morgenstern-Platz 1, 1090 Vienna, Austria and Department of Economics, CentER, TILEC, Tilburg University, PO Box 90153, 5000 LE Tilburg, The Netherlands; E-mail: wieland.mueller@univie.ac.at.

1 Introduction

This paper is concerned with the experimental occurrence of collusion (low quantities) in Cournot environments. The novelty is that we introduce repeated Bayes-Nash Cournot games, where two firms repeatedly, independently, and privately draw their cost in each round, and compare them to environments where firms either have the same costs or have different, but constant and known, costs.

The emergence of (tacit) collusion in oligopolistic environments is of particular interest. For one, collusion in oligopoly typically takes the form of a social dilemma (with Nash equilibrium predictions conflicting with the collective interests of the players, as exemplified by the difference between Cournot equilibrium profits and monopoly profits). Cooperation in social dilemmas is the subject of an enormous literature across the various social sciences. Second, in many countries, the fight against collusion and cartels is at the top of competition authorities' concerns so that work on the determinants of collusion in oligopolistic markets can directly inform public policy.

Cournot competition is a workhorse of industrial organization and it has been extensively studied in the lab. Some of the very first studies in experimental economics concerned themselves with behavior in Cournot environments (see Sauermann and Selten, 1959; Hoggatt, 1959). Many determinants of collusion in Cournot environments have now been explored: the number of firms (Huck et al., 2004); the possibility of pre-play communication (Binger et al., 1990; Waichman, Requate, and Siang, 2014; Fischer and Normann, 2019); the type of feedback information about play (Huck et al., 2000; Davis 2002; Offerman et al., 2002; Altavilla et al. 2006); the matching protocol (Davis et al., 2003); the frequency or duration of interaction (Normann, Requate, and Waichman, 2014; Bigoni, Potters, and Spagnolo, 2019); the use of complete contingent strategies (as opposed to making a choice in every round; see Selten, Mitzkewitz, and Uhlich, 1997); gender effects (Mason, Phillips, and Redington, 1991); the level of the discount rate (Feinberg and Husted, 1993); or the nature of decision-making (Raab and Schipper, 2009).¹

The majority of studies look at symmetric environments, where firms have the same cost functions. Some studies (Fouraker and Siegel, 1963; Mason, Phillips, and Nowel, 1992; Mason and Phillips, 1997; Selten, Mitzkewitz, Uhlich, 1997; Rassenti et al., 2000; Normann, Requate, Waichman, 2014; Fischer and Normann, 2019) introduce asymmetric costs. Overall, there is evidence that asymmetry makes it harder to collude in the lab and observed quantities are typically higher than in symmetric configurations.

Most of this literature looks at environments with complete information. Exceptions include: Fouraker and Siegel (1963), Carlson (1967), Mason and Phillips (1997), which have conditions in which a given player does not have any information about the payoff of the other player(s).² Not all of those studies com-

¹For a meta-study on the determinants of collusion in oligopoly experiments, more generally, see Engel (2007).

²Thus, in those studies, the games are not Bayesian, in the technical sense (Harsanyi, 1967) of having a prior distribution of types commonly known by all players.

pare complete information to incomplete information environments; when they do, they report a tendency for collusion to be harder to achieve in the presence of complete information.

To our knowledge, no paper so far has looked at a (finitely) repeated standard Bayes Cournot game with uncertain costs, an environment which displays both incomplete information and potential changes in asymmetric cost levels (as cost types are drawn anew in every round). In fact, there is a scarcity of articles looking at repeated Bayes-Nash environments in the more general literature about experimental oligopolies. We are only aware of a study by Abbink and Brandts (2005), which speaks to the possibility of collusion in Bayes-Nash *Bertrand* oligopolies. In their experiment, Bertrand firms face a known linear demand curve but they independently and repeatedly draw their unit cost from a common (uniform) distribution under fixed matching for 50 rounds.

A Bayes-Nash Cournot environment is interesting to study for several reasons. First, that game is canonical and, as such, worthy of investigation. Second, under finite repetition, subgame-perfection predicts that players will play the unique, static Bayes-Nash equilibrium in every round and one would want to know whether that prediction will be borne out. In complete information duopolies, Cournot players typically manage, under sufficiently long repetition, to achieve higher payoffs than predicted by the Cournot equilibrium. The previous literature suggests that the presence of (possible) cost asymmetry or incomplete information should complicate the task of subjects in our context but it is simply not known to which extent collusion might be impaired and whether that depends on the magnitude of the asymmetry. Third, outside the lab firms are likely to have private information about their (changing) cost level. Although arguably specific, the Bayes Cournot environment brings a measure of stochasticity to a literature which has mainly focused on very stable (indeed, identically repeated) contexts.

We conduct a series of laboratory experiments with two players under fixed matching and finite repetition. Our treatments include full symmetry and complete information, some cost asymmetry under complete information, and private information about repeatedly drawn costs (the proper Bayes-Nash treatments). Subjects remain matched to the same partner for 60 rounds and face the same, known linear demand curve. In the Bayes-Nash treatments, in every round, costs are drawn to be high or low with equal probability. In a sequence of treatments, we vary the level of asymmetry (i.e. the distance between the high and the low cost).

We uncover the following main findings. Whereas for markets with complete cost symmetry and complete information our data reproduce the known result that duopoly players achieve on average partially collusive outcomes (see e.g. Huck et al., 2004), we find that as soon as any level of cost asymmetry or incomplete information is introduced observed average individual quantities are remarkably close to the static (Bayesian) Nash equilibrium values. Moreover, we do not observe substantial differences in collusion levels among treatments based on the size of the cost asymmetry.

We investigate the adjustment process of decision-making by subjects from one round to the next (‘learning’) in the spirit of Huck, Normann and Oechssler (1999) and Rassenti et al. (2000). We find evidence that in the treatments where either cost asymmetry or incomplete information is present, subjects’ adjustments are more in line with Cournot best-response to the opponent’s previous choice rather than with imitation of it. By contrast, the symmetric, complete-information treatment is the only one where players put less weight on playing a best-response to their opponent’s last round choice and more weight on imitating it (thus, allowing players to find their way towards cooperation by achieving gradual reductions in output).

We conclude that there is something special to the treatment involving two players under symmetry, complete information and finite repetition, which leads players to depart more from myopic optimization. In the other treatments, we find that the static Bayes-Nash equilibrium values are good predictors. In that sense, observed behavior is ‘discontinuous’ as soon as one moves away from complete information and full symmetry.

This finding reinforces the idea that tacit collusion can be achieved in Cournot environments only in very specific circumstances. Remarkably (and setting external validity concerns aside for a moment), this seems to align well with the decisional practice of competition authorities when ruling on the so-called “coordinated effects” (i.e. the possibility of tacit collusion) in merger control. Davies, Olczak and Coles (2011) indeed show that the European Commission concerns itself with collusion threats only in the case of post-merger *symmetric* duopolies.

The rest of this paper is structured as follows. In Section 2, we briefly describe the standard theoretical predictions associated with our Bayes-Nash environment. Section 3 describes our experimental set-up and the various treatments. Section 4 contains our findings and our analysis of subjects’ adaptive behavior. Section 5 concludes.

2 Theory

Consider an incomplete-information Cournot duopoly operating in a market with inverse demand

$$P(Q) = \max\{a - bQ, 0\},$$

where $Q = q_1 + q_2$ is the aggregate quantity in the market. Suppose that firm $i = 1, 2$ has unit costs c_i^H with probability λ_i and c_i^L with probability $1 - \lambda_i$, where $c_i^H \geq c_i^L \geq 0$, and that these costs are privately observed.

Let $q_i^*(c_i^H)$ and $q_i^*(c_i^L)$ denote the quantities produced by firm $i = 1, 2$ in the Bayes-Nash equilib-

Table 1: Experimental Design

Treatment	Info	Parameter Choices	#Subjects	#Markets	#Obs
In all Treatments: $a = 120$ and $b = 1$					
T-30,30-C	C	$c_1 = c_2 = 30$	14	7	840
T-29,31-C	C	$c_1^H = c_2^H = 31, c_1^L = c_2^L = 29$	30	15	1,800
T-29,31-I	I	$\lambda_1 = \lambda_2 = 0.5, c_1^H = c_2^H = 31, c_1^L = c_2^L = 29$	28	14	1,680
T-25,35-I	I	$\lambda_1 = \lambda_2 = 0.5, c_1^H = c_2^H = 35, c_1^L = c_2^L = 25$	28	14	1,680
T-20,40-I	I	$\lambda_1 = \lambda_2 = 0.5, c_1^H = c_2^H = 40, c_1^L = c_2^L = 20$	26	13	1,560
T-20,40-10,50-I	I	$\lambda_1 = \lambda_2 = 0.5, c_1^H = 40, c_1^L = 20, c_2^H = 50, c_2^L = 10$	28	14	1,680

Notes: The letter I (C) in column Info indicates that firms have (In)Complete information about each other's costs. The individual equilibrium quantities for each treatment are indicated in Table 2 in the column labeled "Quantity Predicted."

rium (BNE) depending on own costs. It is routine to show that these quantities are given by:

$$q_i^*(c_i^L) = \frac{1}{6b} (2a - 4c_i^L + 2c_j^L - \lambda_i c_i^H + 2\lambda_j c_j^H + \lambda_i c_i^L - 2\lambda_j c_j^L) \quad (1)$$

$$q_i^*(c_i^H) = \frac{1}{6b} (2a - 3c_i^H - c_i^L + 2c_j^L - \lambda_i c_i^H + 2\lambda_j c_j^H + \lambda_i c_i^L - 2\lambda_j c_j^L), \quad (2)$$

where $i, j = 1, 2$ and $i \neq j$.

In a Cournot duopoly with complete information about (possibly different) costs $c_i \geq 0$ firms choose the following quantities in the Nash equilibrium (just set $c_k^L = c_k^H = c$ for $k = i, j$ in (1) or (2))

$$q_i^* = \frac{1}{3b} (a - 2c_i + c_j), \quad i, j = 1, 2 \text{ and } i \neq j. \quad (3)$$

Provided the BNE of a stage game is unique, in a finitely repeated Bayesian game, the only perfect Bayesian equilibrium is to play the stage-game BNE in every round of the repeated game. While the collusive outcome in a symmetric Cournot duopoly with complete information is clear (each firm produces half the monopoly quantity), in a Cournot duopoly with asymmetric costs and complete information players do not agree on the collusive actions (see Schmalensee (1987) for a theoretical and Fischer and Normann (2019) for a theoretical and experimental investigation of this case). Finally, in a Bayesian Cournot game cooperation/collusion problems are arguably even more severe, due to private information about costs.

3 Experimental Design and Procedures

In the experiment, subjects participated in 60 consecutive rounds. In each round, the inverse demand function was given by $P(Q) = \max\{0, 120 - Q\}$, where $Q = q_1 + q_2$ represents the aggregate quantity in the

market. Participants acted as firms and decided simultaneously on their quantities q_i , $i = 1, 2$. We employed a between-subjects design. Table 1 gives an overview of all treatments. The treatments differ with respect to the distribution of the unit costs of the two firms, c_i , and with respect to the information about the cost structure in the market. In two out of the six treatments (indicated with the letter ‘C’ in the treatment’s name), the costs of both firms in a round were common knowledge. In the other four treatments (with the letter ‘I’ in the treatment’s name), subjects only knew their own costs. More precisely, T-30,30-C is a standard Cournot duopoly in which firms have constant unit costs of 30 each throughout the experiment and know it. In all other treatments, firms have one of two possible unit costs in each round, where in each round the unit costs are randomly and independently assigned with probability 0.5. While in T-29,31-C both firms know their own and the other firm’s unit costs in each round, in all I-treatments each firm knows (a) its own randomly assigned unit cost and (b) the binary distribution of the unit cost of the other firm but not its realization. Note that three of the four I-treatments are ex-ante symmetric, while the fourth, T-20,40-10,50-I, is asymmetric as one firm has the two possible costs level of 20 and 40 and the other 10 and 50, respectively. Finally, note that in all treatments the ex-ante expected costs of firms are equal to 30.

In each round, subjects could choose a non-negative quantity not larger than 120 with the smallest step size being 0.01. Before making their quantity decision, subjects also had the opportunity to simulate different market scenarios with the help of a profit calculator: they could enter two arbitrary quantities, one for themselves and one for their opponent, and were then shown the resulting profit for them.³ After all subjects had submitted their decisions, the computer software cleared the market by quoting the price leading (simulated) demand to equal the entire fictional quantity supplied. Subjects were then informed about the following: the last round’s cost information (own cost in I treatments or both costs in C treatments), the quantity decisions of both firms, and their own profit in this round. This information remained present on the screen when deciding in the next round. Note that no information about the unit cost of the other firm was ever provided in the incomplete-information treatments.

Upon arrival in the lab, participants were given written instructions (see the Appendix for a translated version). Each participant was assigned to a computer and randomly matched with another subject with whom they interacted over the entire experiment. Subjects never learnt with whom they formed a market and it was made sure that communication among subjects was not possible. However, it was common knowledge that the composition of markets formed at the beginning of the experiment remained fixed throughout the whole experiment. The instructions stated that subjects would represent a firm in a market competing with one another firm.

The experiment was programmed and conducted using zTree (Fischbacher 2007) at the Technical

³The profit calculator, provides essentially the same information as commonly used payoff tables, but helps to avoid a possible bias due to limited computational abilities of participants (Huck et al. 2000, p. 42).

University Berlin and Humboldt University Berlin. Participants were students (32% female), mostly from economics, business, natural sciences, or engineering. Altogether, we conducted 77 markets with 154 subjects and collected 9,240 quantity decisions. Each subject participated in one market only.

In the experiment, a fictional currency called ECU (Experimental Currency Unit) was used, with a pre-announced exchange rate of 3000 ECU = 1 EUR. Subjects were made aware of the fact that their profit could become negative in case the market price exceeded their unit costs. For this reason subjects received a starting balance of 7500 ECU and were told that in case of a loss the negative payoff would be offset against their cumulative profits so far in addition to their starting balance. At the end of the experiment, they were paid on the basis of their cumulated earnings over the 60 rounds of play. Sessions took about 60 minutes to complete. The average total earnings per subject was 20.43 EUR.

4 Experimental Results

4.1 Aggregated results

Table 2 provides summary statistics of our experimental results. We provide averages of individual quantities per market (with standard errors of the mean in parentheses) for various time intervals and for each of our treatments separately.⁴ Table 2 also shows the results of two-tailed one-sample *t*-tests of whether the sample mean is equal to the theoretically predicted value. The unit of observation for the tests are market averages of individual quantities. Looking at Table 2, we make a number of observations. First, for treatment T-30,30-C, we find confirmation of the known result that subjects are, on average, (partially) able to collude. For the three time intervals considered, the *t*-test indicates that the observed individual market averages are statistically significantly below the Nash equilibrium.⁵ Second, in all other treatments the observed averages are remarkably close to the Bayes-Nash equilibrium values (and in only very rare cases do we observe that the *t*-test only weakly significantly rejects equality of observed averages with predicted values). This is perhaps most surprising in treatments T-29,31-C and T-29,31-I where subjects know that in each round they have very similar costs. Yet it appears that subjects are unable to collude successfully even though they interact repeatedly over 60 rounds in fixed pairs. Figure 1 in the Appendix shows the distributions (histograms) of averages of individual quantities per market for each treatment separately.

The question is why we observe successful collusion in treatment T-30,30-C but neither in the other complete-information treatment T-29,31-C nor in any of the incomplete-information treatments. We conjecture that except in treatment T-30,30-C subjects just do not know what quantity or quantities to collude on as market conditions change (albeit only slightly in some of the treatments) from round to round. Note, addi-

⁴As mentioned before, collusive outcomes are unclear in all treatments but T-30,30-C. Hence, we do not provide collusion indices (Friedman 1971) as is customary in many papers on market experiments.

⁵Note that the individual collusive quantity is 22.5.

Table 2: Summary statistics

Treatment	Costs	Quantity in the BNE	Average Individual Quantity Observed		
			Rounds 1-30	Rounds 31-60	Rounds 1-60
T-30,30-C	$c = 30$	30	26.72** (1.26)	25.72** (1.24)	26.22** (1.09)
T-29,31-C ^a	$c_1^L = 29, c_2^L = 29$	30.33	28.81* (0.79)	30.27 (0.71)	29.69 (0.66)
	$c_1^L = 29, c_2^H = 31$	31	30.75 (0.78)	30.68 (0.47)	30.70 (0.58)
	$c_1^H = 31, c_2^L = 29$	29	28.15 (0.80)	29.19 (0.54)	28.72 (0.64)
	$c_1^H = 31, c_2^H = 31$	29.67	28.49* (0.64)	29.75 (0.50)	28.99 (0.56)
T-29,31-I	$c^L = 29$	30.5	29.72 (0.85)	29.26 (0.95)	29.49 (0.88)
	$c^H = 31$	29.5	28.04 (0.73)	28.89 (0.92)	28.45 (0.78)
T-25,35-I	$c^L = 25$	32.5	32.22 (0.77)	32.30 (0.53)	32.27 (0.57)
	$c^H = 35$	27.5	27.63 (0.41)	28.25 (0.70)	27.93 (0.49)
T-20,40-I	$c^L = 20$	35	34.52 (0.71)	35.69 (0.77)	35.12 (0.65)
	$c^H = 40$	25	24.96 (0.58)	25.27 (0.78)	25.08 (0.56)
T-20,40-10,50-I	$c_1^L = 20$	35	35.00 (1.46)	36.50 (1.96)	35.64 (1.62)
	$c_1^H = 40$	25	24.18 (1.45)	23.44 (1.17)	23.77 (1.24)
	$c_2^L = 10$	40	37.72 (1.84)	39.16 (2.20)	38.49 (1.88)
	$c_2^H = 50$	20	19.86 (1.04)	17.90* (1.17)	18.93 (1.09)

Notes: This table shows averages of individual quantities per market with standard errors of the mean in parentheses. BNE refers to the Bayesian Nash equilibrium. ^a In treatment T-29,31-C, BNE and observed quantities refer to those of player 1. Test statistics refer to two-tailed one-sample *t*-tests of whether the sample mean is equal to BNE quantities. The unity of observation for the tests are averages of individual quantities per market. The symbols **, * indicate significance at the 5%, 10% level.

tionally, that for asymmetric costs in treatment T-29,31-C the two players disagree about the collusive action (Schmalensee, 1987 and Fischer and Normann, 2019). The stability of the market environment in treatment T-30,30-C seems to enable successful coordination, while this is not the case in all other treatments.

In the next subsection, we shed light on this issue by estimating to what extent behavior in our treatments accords with various well-known learning dynamics.

4.2 Learning dynamics

In view of results in the literature (Huck et al. 1999, 2002, Rassenti et al. 2000) and the inspection of our data, we consider the following learning dynamics.

Best-response dynamics. According to this dynamics, player i chooses some quantity in round $t = 1$ and in round $t \geq 2$ chooses a best response, denoted by r_i^{t-1} , to the other player's quantity in the previous round, q_j^{t-1} . In case of a complete-information Cournot duopoly with linear demand and costs, this dynamic converges to the Nash equilibrium of the one-shot game given in (3); see Theocharis (1960). In the incomplete-information games considered in this study, the best response dynamics can be shown (own simulations) to converge on average to the Bayes-Nash equilibrium given in (1) and (2), where individual quantities oscillate in the interval

- $[q_i^*(c_i^k) - \Delta, q_i^*(c_i^k) + \Delta]$ with $k \in \{L, H\}$, $q_i^*(c_i^k)$ given in (1) and (2), and $\Delta = (q_i^*(c_i^L) - q_i^*(c_i^H))/2$ in case of treatments T-29,31-I, T-25,35-I and T-20,40-I;
- $[q_i^*(c_i^k) - \Delta_i, q_i^*(c_i^k) + \Delta_i]$ with $k \in \{L, H\}$, $q_i^*(c_i^k)$ given in (1) and (2), and $\Delta_i = q_i^*(c_i^H)/3$, $i = 1, 2$ in case of treatment T-20,40-10,50-I.

Fictitious-play dynamics. According to this dynamics, player i chooses some quantity in round $t = 1$ and in round $t \geq 2$ chooses a best response, denoted by f_i^{t-1} , to the average of the other player's quantities in all previous rounds, $\frac{1}{t-1} \sum_{k=1}^{t-1} q_j^k$. This dynamics can be shown to converge (Nachbar (1990) and own simulations) to the (Bayesian) Nash equilibrium quantities given in (1), (2) and (3).

Imitate the other. According to this dynamics, player i chooses some quantity in round $t = 1$ and in round $t \geq 2$ either chooses what the other player chose in the previous round (the version we call “imitate last round” and denote by il_i^{t-1}) or the average of the other player's quantities in all previous rounds (the version we call “imitate fictitious” and denote by if_i^{t-1}). Note that the first version is simply the learning rule “imitate the average,” also used in Huck et al. (1999, 2002), which in our duopoly context simply means “imitate the other firm.”⁶ The “imitate last round” dynamic does not converge and perpetually jumps between the initial choices of the two players. The “imitate fictitious” dynamics converge in three rounds to the average of the two players' initial choices in all treatments.

⁶Note that the rule “imitate the average” should not be confused with the imitation rule analysed by Vega-Redondo (1997), in which a player, when given the opportunity to revise its choice, imitates the firm with the highest payoff in the last round or chooses randomly with some positive probability.

Several remarks are in order. First and most importantly, the applicability of these dynamics to all of our treatments can be challenged. For instance, the strict best-response dynamics that only takes into account the quantity chosen in the previous round can be questioned in case of treatment T-29,31-C as well as in the incomplete-information treatments: in treatment T-29,31-C, players face different combinations of c_i and c_j in every round and know about it, and in case of the incomplete-information treatments, players should best respond to the expected quantity given the two possible cost levels. Note, though, that players in the incomplete-information treatments only observe the quantity chosen, but not the type of the other player in the previous round. Second, questions also arise as to whether the fictitious-play dynamics should be applied in all treatments. For example, in treatment T-29,31-C in which players know their own and the other player's cost when making a decision, one might argue that players should only use relevant previous rounds to form their fictitious-play beliefs, that is, those rounds in which both players had the same costs as in the current round.⁷

Despite these qualifications, we run estimations based on all the dynamics above for all treatments as it is an open question which rounds of earlier play subjects take into account (if any) when deciding about the choice in the current round. Moreover, we invoke all the dynamics for all treatments to compare their performance on equal terms across all treatments.

We estimate two different models using two ways of constructing regressors of the estimation equations. We explain these in turn. The first model we estimate (see also Huck et al., 1999, 2002; Rassenti et al., 2000) is

$$q_i^t - q_i^{t-1} = \beta_0 + \beta_1(r_i^{t-1} - q_i^{t-1}) + \beta_2(il_i^{t-1} - q_i^{t-1}), \quad (4)$$

where r_i^{t-1} denotes subject i 's best response to the other firms quantity in $t-1$ and il_i^{t-1} denotes the quantity of the other firm's quantity in $t-1$. The second model we estimate is

$$q_i^t - q_i^{t-1} = \gamma_0 + \gamma_1(f_i^{t-1} - q_i^{t-1}) + \gamma_2(if_i^{t-1} - q_i^{t-1}), \quad (5)$$

where f_i^{t-1} is the best reply against fictitious play beliefs, and if_i^{t-1} is the average of the quantities of i 's rival in all previous rounds. We do not estimate an equation in which we combine the regressors of equations (4) and (5), because the terms $(r_i^{t-1} - q_i^{t-1})$ and $(f_i^{t-1} - q_i^{t-1})$ (as well as the terms $(il_i^{t-1} - q_i^{t-1})$ and $(if_i^{t-1} - q_i^{t-1})$) are highly correlated with each other.

Clearly, the size of the estimated coefficients $\hat{\beta}_1$ and $\hat{\beta}_2$ in model (4) and $\hat{\gamma}_1$ and $\hat{\gamma}_2$ in model (5), respectively, indicate whether subjects tend to play a best response to the last or to the average of previous rounds or tend to imitate what the other firm did in the last round or the previous rounds on average.

We estimate models (4) and (5) using two different ways of constructing regressors. In the regres-

⁷Below we explain how we account for these qualifications.

sions labeled “previous rounds” we compute r_i^{t-1} and il_i^{t-1} for model (4) using q_{-i}^{t-1} , where, with abuse of notation, $t - 1$ refers to the “truly” previous round; and compute f_i^{t-1} and if_i^{t-1} for model (5) by using the average of the quantities of i ’s rival in all previous rounds. In the regressions labeled “previous relevant rounds” we use q_{-i}^{t-1} , where $t - 1$ now refers to the most recent “relevant” round. More precisely, say that in the current round $c_i = c_j = 29$. Then $t - 1$ refers to the most recent round in which firms had the same cost information as in the current round. To compute f_i^{t-1} and if_i^{t-1} for model (5), we use the average of all quantity choices by i ’s rival in all previous “relevant” rounds. More precisely, say that in the current round $c_i = c_j = 29$. Then we used the average of all quantities chosen by i ’s rival in all previous rounds in which $c_i = c_j = 29$.

In treatment T-30,30-C, the costs of players remain the same across all rounds. Hence, the two ways of constructing regressors are the same and so we just estimate models (4) and (5) once. In all other treatments, cost assignments are random across rounds. Hence, we estimate models (4) and (5) twice, once using “previous rounds” data and once using “previous relevant rounds” data. We estimated the models by a mixed-effects, multilevel panel data model where subjects are nested in markets, controlling for heteroskedasticity. The results are given in Table 3.

Inspection of Table 3 reveals that treatment T-30,30-C is the only one for which we find the estimated constants to be significantly smaller than zero. This indicates a downward time trend of average chosen quantities. Second, the estimated coefficients of the main regressors are positive in all treatments but clearly below 1, which means that adaptations do not accord fully with any of the dynamics specified. Third, for treatment T-30,30-C we find that $\hat{\beta}_1$ is significantly smaller than $\hat{\beta}_2$ in model (4), indicating that subjects on average choose more in accordance with imitation of the other firm’s last-round choice than best responding to it. In model (5) of T-30,30-C, we find that $\hat{\gamma}_1$ is also smaller than $\hat{\gamma}_2$. However, the difference is not statistically significant. Fourth, and most importantly, in all other treatments we observe that the estimated coefficient $\hat{\beta}_1$ ($\hat{\gamma}_1$) is significantly larger than the estimated coefficient $\hat{\beta}_2$ ($\hat{\gamma}_2$), indicating that subjects in the treatments with either cost asymmetry or incomplete information choose on average more in accordance with best-response behavior than with imitation.⁸ In fact, neglecting two extreme cases, the mean of the ratios $\hat{\beta}_1/\hat{\beta}_2$ and $\hat{\gamma}_1/\hat{\gamma}_2$ is 2.5 with a standard deviation of 1.4 in all treatments other than T-30,30-C. Finally, note that the *Log LL* values for the two models estimated for the same way of constructing regressors are usually very close to each other, which indicates that none of the models is clearly favored in accounting for the way subjects adapt and learn over time. This is arguably not surprising given the earlier observation regarding the correlation between regressors in models (4) and (5).⁹

⁸There is one exception, namely, model (5) in treatment T-25,31-I when estimated using previous round data.

⁹We repeated the analysis reported in Table 3 (that uses all data), for the data of the first half of the experiment only. We did so as adjustments might be particularly pronounced at the beginning of the experiment. We find very similar results to those shown in Table 3 with the one exception that for treatment T-30,30-C the coefficients β_1 than β_2 in model (4) are not statistically different from each other anymore. Hence, also for the first half of the data we find that subjects’ adjustments are on average more in line

Table 3: Results of adjustment regressions

	Treatment T-30,30-C				Treatment T-29,31-C			
					“Previous rounds”		“Previous relevant rounds”	
	Model (4)	Model (5)	Model (4)	Model (5)	Model (4)	Model (5)	Model (4)	Model (5)
β_1 / γ_1	0.204*** (0.067)	0.211*** (0.054)	0.490*** (0.066)	0.570*** (0.089)	0.520*** (0.026)	0.535*** (0.038)		
β_2 / γ_2	0.399*** (0.050)	0.276*** (0.077)	0.227*** (0.039)	0.152*** (0.026)	0.254*** (0.029)	0.240*** (0.024)		
β_0 / γ_0	-1.126*** (0.368)	-1.283*** (0.336)	-0.283 (0.395)	-0.416 (0.457)	-0.144 (0.424)	-0.171 (0.425)		
N	826	826	1,770	1,770	1,680	1,680		
Log LL	-2082	-2116	-5061	-5053	-4637	-4636		
p -value of H_0	0.044	0.410	< 0.001	< 0.001	< 0.001	< 0.001		
$\hat{\beta}_1 = \hat{\beta}_2 / \hat{\gamma}_1 = \hat{\gamma}_2$								

	Treatment T-29,31-I				Treatment T-25,35-I			
	“Previous rounds”		“Previous relevant rounds”		“Previous rounds”		“Previous relevant rounds”	
	Model (4)	Model (5)	Model (4)	Model (5)	Model (4)	Model (5)	Model (4)	Model (5)
β_1 / γ_1	0.500*** (0.045)	0.509*** (0.057)	0.582*** (0.043)	0.537*** (0.065)	0.547*** (0.048)	0.544*** (0.059)	0.491*** (0.029)	0.496*** (0.035)
β_2 / γ_2	0.296*** (0.044)	0.289*** (0.069)	0.291*** (0.026)	0.336*** (0.053)	0.386*** (0.036)	0.391*** (0.062)	0.294*** (0.027)	0.288*** (0.052)
β_0 / γ_0	-0.816 (0.596)	-0.827 (0.628)	-0.919 (0.728)	-0.838 (0.701)	0.154 (0.388)	0.233 (0.377)	0.173 (0.343)	0.195 (0.347)
N	1,652	1,652	1,624	1,624	1,652	1,652	1,624	1,624
Log LL	-4987	-4989	-4897	-4895	-4988	-4998	-4865	-4867
p -value of H_0	< 0.001	0.021	< 0.001	0.072	0.023	0.165	< 0.001	0.010
$\hat{\beta}_1 = \hat{\beta}_2 / \hat{\gamma}_1 = \hat{\gamma}_2$								

	Treatment T-20,40-I				Treatment T-20,40-10,50-I			
	“Previous rounds”		“Previous relevant rounds”		“Previous rounds”		“Previous relevant rounds”	
	Model (4)	Model (5)	Model (4)	Model (5)	Model (4)	Model (5)	Model (4)	Model (5)
β_1 / γ_1	0.710*** (0.042)	0.824*** (0.075)	0.402*** (0.049)	0.505*** (0.062)	0.742*** (0.029)	0.935*** (0.093)	0.389*** (0.094)	0.530*** (0.126)
β_2 / γ_2	0.301*** (0.035)	0.142** (0.060)	0.156*** (0.033)	0.079* (0.045)	0.315*** (0.024)	0.056 (0.098)	0.148*** (0.045)	0.046 (0.051)
β_0 / γ_0	0.133 (0.393)	0.106 (0.469)	0.138 (0.235)	0.132 (0.293)	-0.797 (0.914)	-1.159 (1.065)	-0.361 (0.488)	-0.593 (0.602)
N	1,534	1,534	1,508	1,508	1,652	1,652	1,624	1,624
Log LL	-4728	-4698	-4424	-4404	-5835	-5775	-5527	-5505
p -value of H_0	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001	< 0.001
$\hat{\beta}_1 = \hat{\beta}_2 / \hat{\gamma}_1 = \hat{\gamma}_2$								

Notes: This table shows the results of the adjustment regressions using all data. For a description of the specifications labeled “previous rounds” and “previous relevant rounds,” see page 10 of the main text. In treatment T-30,30-C, the costs of players remain the same across all rounds. Hence, the two ways of constructing regressors are the same and so we just estimate models (4) and (5) once. In all other treatments, cost assignments are random across rounds. Hence, we estimate models (4) and (5) twice, once using “previous rounds” data and once using “previous relevant rounds” data. The β / γ coefficients refer to model (4) / (5). The symbols ***, **, * indicate significance at the 1%, 5%, 10% level.

We checked these results also at the individual level. That is, for each subject of a treatment separately we estimated models (4) and (5) for both ways of construction regressors, controlling for autocorrelation and heteroskedasticity. The results are presented in Table 4 in the appendix. The entries in this table indicate the share of subjects per treatment for which the hypothesis listed in the second column of Table 4 is rejected. We make several observations. First, in T-30,30-C the share of subjects for which the hypothesis of $H_0: \beta_0 \geq 0$ ($\gamma_0 \geq 0$) is rejected in favor of $H_1: \beta_0 < 0$ ($\gamma_0 < 0$) is clearly positive, while the share of subjects for which $H_0: \beta_0 \leq 0$ ($\gamma_0 \leq 0$) is rejected in favor of $H_1: \beta_0 > 0$ ($\gamma_0 > 0$) is zero. Note that the corresponding numbers in all other treatments are usually more similar to each other. This indicates that also at the individual level, in treatment T-30,30-C there is on average more of a downward trend in chosen quantities than an upward trend. Second, in all treatments but T-30,30-C the share of subjects for which the hypothesis of $H_0: \beta_1 \leq \beta_2$ ($\gamma_1 \leq \gamma_2$) is rejected in favor of $H_1: \beta_1 > \beta_2$ ($\gamma_1 > \gamma_2$) is always much larger than the share of subjects for which $H_0: \beta_1 \geq \beta_2$ ($\gamma_1 \geq \gamma_2$) is rejected in favor of $H_1: \beta_1 < \beta_2$ ($\gamma_1 < \gamma_2$). Note that the corresponding numbers in Treatment T-30,30-C are the same. This indicates that also at the individual level, in all treatments but T-30,30-C subjects' adjustments are on average more in line with best-response behavior than with imitation.

4.3 Additional evidence

The analysis of the recorded simulations conducted by subjects prior to the actual quantity choices confirms the observed difference in the subjects' decision approach used in treatment T-30,30-C compared to the treatments with either cost asymmetry or incomplete information.¹⁰ For example, subjects in treatment T-30,30-C used the profit calculator least often, and also the share of actual quantity choices tried out in the simulations were at the minimum in treatment T-30,30-C (see Table 5 in the Appendix for more detailed results from the simulation data analysis). The observed "discontinuity" in behavior is also confirmed by subjects' answers in the post-experimental questionnaire regarding the question of how they came to their decisions in the experiment. For example, in T-30,30-C, the word "collusion" or a description of an attempt to achieve collusion was mentioned by 71% of the subjects, while the corresponding share in all other treatments is not higher than 39%.

with best-response behavior than with imitation in all treatments but treatment T-30,30-C.

¹⁰Recall that according to our experimental design, before making their quantity decisions, subjects had the opportunity to simulate different market scenarios with the help of a profit calculator. More precisely, they could try different pairs of quantities (own and of the opponent) and were then shown the resulting profit for themselves.

5 Conclusions

We report on Cournot duopoly market experiments with a relatively high number of repetitions and fixed matching. We run treatments that include markets with (a) complete cost symmetry and complete information, (b) slight cost asymmetry and complete information, and (c) varying cost asymmetries and incomplete information.

The main result can be interpreted as a “discontinuity” in behavior: While for markets with complete cost symmetry and complete information our data confirm the known result that duopoly players achieve on average partially collusive outcomes, we find that as soon as any level of cost asymmetry or incomplete information is introduced observed average individual quantities are remarkably close to the static (Bayesian) Nash equilibrium values. This is so despite repeated and fixed matching over the course of 60 rounds.

The results of various regressions analyzing players’ adjustment behavior over time provide an explanation of this main result. We find significantly more adjustments in line with best-response behavior than with imitation in all but the treatment with complete symmetry and information. This provides an explanation of our results as simulations show that best-response dynamics do converge (on average) to static (Bayesian) Nash equilibrium quantities.

In their duopoly treatment (as well as the ones with 3 or 4 firms), Abbink and Brandt (2005) found that prices were systematically below the Bayes-Nash values, that is, observed play was more competitive than Bayes-Nash equilibrium predictions. We find, on the contrary, that, in our incomplete information treatments, observed average quantities are in line with Bayes-Nash equilibrium predictions. This may yet again point to a fundamental difference between experimental Bertrand and Cournot environments (and more generally, games of strategic substitutes vs. games of strategic complements, see e.g. Potters and Suetens, 2009; Mermer, Müller and Suetens, 2021). Note, however, that, in contrast to the evidence relating to complete-information, symmetric contexts (Suetens and Potters, 2007), in Bayes-Nash environments Bertrand appears to lead to more competitive outcomes than Cournot.

The adjustment dynamics we explore in this paper have so far mainly been applied to data of symmetric and complete-information markets. We find that none of the models we specify (models (4) or (5)) or the way regressors are constructed (using strictly previous rounds or previous “relevant” rounds) is clearly favored in accounting for the way subjects adapt and learn over time. The appropriateness of applying some of these dynamics in markets with incomplete-information is debatable. Future theoretical and econometric work should, hence, probe whether alternative specifications can better account for players’ adaptations over time.

References

- [1] Abbink, K. and J. Brandts (2005): Price competition under cost uncertainty: A laboratory analysis, *Economic Inquiry* 43, 636–648.
- [2] Altavilla, C., L. Luini and P. Sbriglia (2006): Social learning in market games, *Journal of Economic Behavior & Organization* 61, 632–652.
- [3] Bigoni, M., J. Potters and G. Spagnolo (2019): Frequency of interaction, communication and collusion: an experiment, *Economic Theory* 68, 827–844.
- [4] Binger, B.R., E. Hoffman, G.D. Libecap and K. Shachat (1990): An experimetric study of the Cournot theory of firm behavior, *University of Arizona* working paper.
- [5] Carlson, J.A. (1967): The stability of an experimental market with a supply-response lag, *Southern Economic Journal* 33, 305–321.
- [6] Davis, D.D. (2002): Strategic interactions, market information and predicting the effects of mergers in differentiated product markets, *International Journal of Industrial Organization* 20, 1277–1312.
- [7] Davis, D.D., R.J. Reilly and B.J. Wilson (2003): Cost structures and Nash play in repeated Cournot games, *Experimental Economics* 6, 209–226.
- [8] Davis, S., M. Olczak and H. Coles (2011): Tacit collusion, firm asymmetries and numbers: evidence from EC merger cases, *International Journal of Industrial Organization* 29, 221–231.
- [9] Engel, C. (2007): How much collusion? A meta-analysis of oligopoly experiments. *Journal of Competition Law & Economics* 3(4), 491–549.
- [10] Feinberg, R.M. and T.A. Husted (1993): An experimental test of discount-rate effects on collusive behaviour in duopoly markets, *Journal of Industrial Economics* 41, 153–160.
- [11] Fischbacher, U. (2007): z-Tree: Zurich toolbox for ready-made economic experiments, *Experimental Economics* 10, 171–178.
- [12] Fischer, C. and H.-T. Normann (2019): Collusion and bargaining in asymmetric Cournot duopoly—An experiment, *European Economic Review* 111, 360–379.
- [13] Fouraker, L.E. and S. Siegel (1963): *Bargaining behavior*. New York: McGraw-Hill.
- [14] Friedman, J. W. (1971): A non-cooperative equilibrium for supergames, *The Review of Economic Studies* 38, 1–12.
- [15] Harsanyi, J.C. (1967): Games with incomplete information played by “Bayesian” players, I—III Part I. The basic model, *Management Science* 14, 159–182.
- [16] Hoggatt, A.C. (1959): An experimental business game, *Behavioral Science* 4, 192–203.
- [17] Huck, S., H.-T. Normann and J. Oechssler (1999): Learning in Cournot Oligopoly—An Experiment, *The Economic Journal* 109, C80–C95.


- [18] Huck, S., H.-T. Normann and J. Oechssler (2000): Does information about competitors' actions increase or decrease competition in experimental oligopoly markets?, *International Journal of Industrial Organization* 18, 39–57.
- [19] Huck, S., H.-T. Normann and J. Oechssler (2002): Stability of the Cournot process – experimental evidence, *International Journal of Game Theory* 31, 123–136.
- [20] Huck, S., H.-T. Normann and J. Oechssler (2004): Two are few and four are many: number effects in experimental oligopolies, *Journal of Economic Behavior & Organization* 53, 435–446.
- [21] Mason, C. F., O. R. Phillips and C. Nowell (1992): Duopoly behavior in asymmetric markets: An experimental evaluation, *Review of Economics and Statistics* 74(4), 662–670.
- [22] Mason, C. F., O. R. Phillips and D.B. Redington (1991): The role of gender in a non-cooperative game, *Journal of Economic Behavior & Organization* 15, 215–235.
- [23] Mason, C. F. and O. R. Phillips (1997): Information and cost asymmetry in experimental duopoly markets, *Review of Economics and Statistics* 79(2), 290–299.
- [24] Mermer, A. G., W. Müller and S. Suetens (2021): Cooperation in infinitely repeated games of strategic complements and substitutes. *Journal of Economic Behavior & Organization* 188, 1191–1205.
- [25] Nachbar, J.H. (1990): “Evolutionary” selection dynamics in games: Convergence and limit properties, *International Journal of Game Theory* 19, 59–89.
- [26] Normann, H.-T., T. Requate and I. Waichman (2014): Do short-term laboratory experiments provide valid descriptions of long-term economic interactions? A study of Cournot markets, *Experimental Economics* 17, 371–390.
- [27] Potters, J. and S. Suetens, S. (2009): Cooperation in experimental games of strategic complements and substitutes. *Review of Economic Studies* 76(3), 1125–1147.
- [28] Offerman, T., J. Potters and J. Sonnemans (2002): Imitation and belief learning in an oligopoly experiment, *Review of Economic Studies* 69, 973–997.
- [29] Raab, P. and B. C. Schipper (2009): Cournot competition between teams: An experimental study, *Journal of Economic Behavior & Organization* 72, 691–702.
- [30] Rassenti, S., S.S. Reynolds, V.L. Smith, and F. Szidarovszky (2000): Adaptation and convergence of behavior in repeated experimental Cournot games, *Journal of Economic Behavior & Organization* 41, 117–146.
- [31] Requate, T. and I. Waichman (2011): “A profit table or a profit calculator?” A note on the design of Cournot oligopoly experiments, *Experimental Economics* 14, 36–46.
- [32] Sauerman, H. and R. Selten (1959): Ein Oligopolexperiment, *Zeitschrift für die gesamte Staatswissenschaft* 115, 427–471.
- [33] Schmalensee, R. (1987): Competitive advantage and collusive equilibria, *International Journal of*

Industrial Organization 5, 351–367.

- [34] Selten, R., M. Mitzekewitz and G.R. Uhlich (1997): Duopoly strategies programmed by experienced players, *Econometrica* 65, 517–555.
- [35] Suetens, S. and J. Potters (2007): Bertrand colludes more than Cournot. *Experimental Economics* 10(1), 71–77.
- [36] Theocharis, R.D. (1960): On the Stability of the Cournot Solution on the Oligopoly Problem, *Review of Economic Studies* 27, 133–134.
- [37] Vega-Redondo, F. (1997): The evolution of Walrasian behavior, *Econometrica* 65, 375–384.
- [38] Waichman, I., T. Requate and C.N.K. Siang (2014): Communication in Cournot competition: an experimental study, *Journal of Economic Psychology* 42, 1–16.

ONLINE APPENDIX
(Not for publication)

Figure 1: Average observed quantities per market



Notes: The panels in this figure show histograms of observed average quantities per market, using the data of all rounds.

Table 4: Summary of hypothesis tests for individual adjustment regressions

	H_0	H_1	T-30,30-C	T-29,31-C	T-29,31-I	T-25,35-I	T-20,40-I	T-20,40-10,50-I
Percentage of subjects for which H_0 is rejected at the 5% level								
“Previous rounds” data, Model (4)	$\beta_0 \geq 0$	$\beta_0 < 0$	23.08	13.33	21.43	17.86	11.54	35.71
	$\beta_0 \leq 0$	$\beta_0 > 0$	0.00	13.33	14.29	17.86	11.54	17.86
	$\beta_1 \geq \beta_2$	$\beta_1 < \beta_2$	23.08	10.00	7.14	0.00	3.85	10.71
	$\beta_1 \leq \beta_2$	$\beta_1 > \beta_2$	23.08	30.00	32.14	50.00	84.62	67.86
“Previous rounds” data, Model (5)	$\gamma_0 \geq 0$	$\gamma_0 < 0$	38.46	10.00	21.43	17.86	26.92	42.86
	$\gamma_0 \leq 0$	$\gamma_0 > 0$	0.00	20.00	17.86	17.86	19.23	10.71
	$\gamma_1 \geq \gamma_2$	$\gamma_1 < \gamma_2$	15.38	20.00	7.14	3.57	0.00	10.71
	$\gamma_1 \leq \gamma_2$	$\gamma_1 > \gamma_2$	15.38	20.00	35.71	25.00	69.23	71.43
“Previous relevant rounds” data, Model (4)	$\beta_0 \geq 0$	$\beta_0 < 0$	23.08	16.67	21.43	7.14	7.69	14.29
	$\beta_0 \leq 0$	$\beta_0 > 0$	0.00	26.67	14.29	14.29	11.54	17.86
	$\beta_1 \geq \beta_2$	$\beta_1 < \beta_2$	23.08	0.00	0.00	3.57	0.00	10.71
	$\beta_1 \leq \beta_2$	$\beta_1 > \beta_2$	23.08	46.67	35.71	28.57	23.08	14.29
“Previous relevant rounds” data, Model (5)	$\gamma_0 \geq 0$	$\gamma_0 < 0$	38.46	16.67	25.00	10.71	15.38	25.00
	$\gamma_0 \leq 0$	$\gamma_0 > 0$	0.00	26.67	14.29	17.86	15.38	28.57
	$\gamma_1 \geq \gamma_2$	$\gamma_1 < \gamma_2$	15.38	0.00	10.71	7.14	3.85	10.71
	$\gamma_1 \leq \gamma_2$	$\gamma_1 > \gamma_2$	15.38	46.67	32.14	21.43	38.46	39.29

Notes: This table shows the results of adjustment regressions at the individual level, using all data. In treatment T-30,30-C one subject had to be excluded as the dependent variable was constant and zero.

Table 5: Statistics from choice simulations prior to actual choices

Treatment	# simulations per subject & round (average)	actual quantity choice was one of simulated quantities (average)	actual quantity choice was last simulated own quantity (relative frequency)	actual quantity choice was last simulated other quantity (relative frequency)	first simulated own quantity was other quantity last round (relative frequency)	first simulated other quantity was own quantity last round (relative frequency)
T-30,30-C	1.46	12.21	40%	41%	35%	49%
T-29,31-C	3.10	28.87	40%	26%	21%	30%
T-29,31-I	4.17	31.96	42%	20%	27%	33%
T-25,35-I	3.23	27.07	49%	19%	23%	23%
T-20,40-I	3.33	29.89	42%	22%	22%	24%
T-20,40-10,50-I	1.84	16.52	36%	12%	15%	13%

Notes: This table shows statistics from subjects' choice simulations prior to actual choices.

Refer to Table 5. The analysis of the recorded simulation data provides additional evidence for the diverse decision approaches used by the subjects in the different treatments. Recall that according to our experimental design, before making their quantity decisions, subjects had the opportunity to simulate different market scenarios with the help of a profit calculator. More precisely, they could try different pairs of quantities (one for themselves and one for the opponent) and were then shown the implied profit for them. The analysis of the simulation data reveals that the decision approach taken in treatment T-30,30-C clearly differs from the approach used in the treatments with either cost asymmetry or incomplete information. For example, subjects in T-30,30-C used the profit calculator least often. The average number of simulations per subject and round is with 1.46 the lowest compared to all other treatments. The actual quantity decision in treatment T-30,30-C was also based least often on the simulation results. In treatment T-30,30-C, in about 12% of the 60 rounds, the subjects' actual quantity choice was equal to one of their own simulated quantities, whereas this number rose up to 28.87% in T-2930-C and 31.96% in T-2939-I, respectively. In T-30,30-C, in 73% of all rounds the actual quantity equals the opponent's quantity chosen in the previous round, whereas in all other treatments this number is not higher than 26%. Finally, in T-30,30-C compared to all other treatments, the first simulated opponent's (own) quantity equals the own (opponent's) actual quantity in the previous round 49% (35%) of all simulations. In all other treatments, those numbers are clearly lower and quite similar in size (see Table 5.)¹¹

¹¹The simulation data are available upon request.

Below we reproduce the translated version of the instructions. For the original instructions (in German), please contact one of the authors. The variants in the instructions for the different treatments are indicated.

Instructions

Please read these instructions carefully. If there is anything you do not understand, please indicate this by raising your hand. We will then answer your questions privately.

In this experiment, you will make decisions repeatedly. In this process you can earn money. How much money you earn depends on your decisions, those of another participant and random moves. The instructions use the fictitious money unit ECU (Experimental Currency Unit). At the end of the experiment, your payouts are converted into euros (see below).

Your anonymity towards us as well as towards the other participants will be preserved.

In this experiment you represent a company that produces and sells one and the same product together with another company on one market. You remain assigned to the same other participant throughout the experiment. All companies always have only one decision to make, namely which quantities they want to produce.

The production costs per unit of your and the other company are determined as follows:

[T-30,30-C]:

- The production costs per unit of your company are 30 ECU.
- The production costs per unit of the other company in your market are 30 ECU.

Afterwards you and the other company decide simultaneously on your quantity.

[T-29,31-C]:

- Regarding the production costs per unit of your company in one round the following holds: with a probability of 50% they will be 31 ECU and with a probability of 50% they will be 29 ECU.
- Regarding the production costs per unit of the other company in one round the following holds: with a probability of 50% they will be 31 ECU and with a probability of 50% they will be 29 ECU.

The production costs per unit of your company and the other company will be chosen independently of each other. Before you decide on the quantity of your company in a round, the production costs per unit of your and the other company are randomly determined and reported to you. Afterwards you and the other company decide simultaneously on your quantity.

[T-29,31-I]:

- Regarding the production costs per unit of your company in one round the following holds: with a probability of 50% they will be 31 ECU and with a probability of 50% they will be 29 ECU.
- Regarding the production costs per unit of the other company in one round the following holds: with a probability of 50% they will be 31 ECU and with a probability of 50% they will be 29 ECU.

The production costs per unit of your company and the other company will be chosen independently of each other. Before you decide on the quantity of your company in a round, the production costs per unit of your company are randomly determined and reported to you. Similarly the production costs per unit for the other company are determined randomly. Each company learns only its own production costs per unit, but not those of the other company. Afterwards you and the other company decide simultaneously on your own quantity.

[T-25,35-I]:

- Regarding the production costs per unit of your company in one round the following holds: with a probability of 50% they will be 35 ECU and with a probability of 50% they will be 25 ECU.
- Regarding the production costs per unit of the other company in one round the following holds: with a probability of 50% they will be 35 ECU and with a probability of 50% they will be 25 ECU.

The production costs per unit of your company and the other company will be chosen independently of each other. Before you decide on the quantity of your company in a round, the production costs per unit of your company are randomly determined and reported to you. Similarly the production costs per unit for the other company are determined randomly. Each company learns only its own production costs per unit, but not those of the other company. Afterwards you and the other company decide simultaneously on your own quantity.

[T-20,40-I]:

- Regarding the production costs per unit of your company in one round the following holds: with a probability of 50% they will be 40 ECU and with a probability of 50% they will be 20 ECU.
- Regarding the production costs per unit of the other company in one round the following holds: with a probability of 50% they will be 40 ECU and with a probability of 50% they will be 20 ECU.

The production costs per unit of your company and the other company will be chosen independently of each other. Before you decide on the quantity of your company in a round, the production costs per unit of your company are randomly determined and reported to you. Similarly the production costs per unit for the other company are determined randomly. Each company learns only its own production costs per unit, but not those of the other company. Afterwards you and the other company decide simultaneously on your own quantity.

[T-20,40-10,50-I]:

- Regarding the production costs per unit of your company in one round the following holds: with a probability of 50% they will be 40 ECU [**50 ECU**] and with a probability of 50% they will be 20 ECU [**10 ECU**].
- Regarding the production costs per unit of the other company in one round the following holds: with a probability of 50% they will be 50 ECU [**40 ECU**] and with a probability of 50% they will be 10 ECU [**20 ECU**].

The production costs per unit of your company and the other company will be chosen independently of each other. Before you decide on the quantity of your company in a round, the production costs per unit of your company are randomly determined and reported to you. Similarly the production costs per unit for the other company are determined randomly. Each company learns only its own production costs per unit, but not those of the other company. Afterwards you and the other company decide simultaneously on your own quantity.

The market price (which can be between 120 ECU and 0 ECU) depends on the total quantity offered by your company and the other company. The following important rule applies: the higher the total quantity of both companies, the lower the price that will be on the market. Moreover, above a certain total quantity, the price becomes zero. More precisely, the price per unit is determined in each round as follows:

$$\text{Price} = 120 - \text{quantity of your company} - \text{quantity of the other company}$$

That means, that in each round the price is equal to the difference between 120 and the total quantity offered by your and the other company. Furthermore, if the total quantity offered by your company and the other company is greater than or equal to 120, the market price is zero.

Your profit per unit in a round is the difference between the market price and your production cost per unit in that round. Note that you make a loss if the market price is less than your unit costs. Your profit in each round is thus equal to the profit per unit times the quantity you chose.

In each round, the quantities of the two companies are recorded, the corresponding price is determined and the respective profits are calculated.

From the second round on, you will be told in each round the quantity of the other company and your own profit of the previous round. For your information, you will be shown your production costs per unit in the previous round and your own quantity in the previous round [**Complete Information-treatments**]: as well as the production costs per unit of the other company. [**Incomplete Information-treatments**]: However, you will not see the production costs per unit of the other company in the previous round.

Before making your choice, you can also simulate your decisions. You can do this on the left side of the decision screen. Here you simply enter any quantity of your own and any quantity of the other company into the two fields and then press the “Compute”-button. In the upper left corner of the screen, you can then see what profit would result for you in that case.

When you have decided on a quantity, enter it in the field on the right side of the screen and press the “OK”-button. Any number between 0 and 120 with two digits after the decimal point can be chosen as a quantity.

The experiment consists of 60 rounds.

Your total payment is the sum of your payments per round. At the end of the experiment, your payments will be converted to Euros, where 3000 ECU = 1€. At the beginning of the experiment, you will receive a (one-time) initial endowment of 7500 ECU.

If you make a loss in a round, it will be deducted from your previous profit (or from your initial endowment).

If there is anything you do not understand, please indicate this by raising your hand. We will then answer your questions privately.