Note

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

Cédric Argenton [a,1], Radosveta Ivanova-Stenzel [b,2], Wieland Müller [c,a,*,3]

[a] *CentER & TILEC, Tilburg University, The Netherlands*
[b] *Technical University Berlin, Germany*
[c] *VCEE, University of Vienna, Austria*

ARTICLE INFO

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.

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 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 Sauerman 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 et al., 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 et al., 2014; Bigoni et al., 2019); the use of complete contingent strategies (as opposed to making a choice in every round; see Selten et al., 1997); gender effects (Mason et al., 1991); the level of the discount rate (Feinberg and Husted, 1993); or the nature of decision-making (Raab and Schipper, 2009).[4]

The majority of studies look at symmetric environments, where firms have the same cost functions. Some studies (Fouraker and Siegel, 1963; Mason et al., 1992; Mason and Phillips, 1997; Selten et al., 1997; Rassenti et al., 2000; Normann et al., 2014; Fischer and Normann, 2019) introduce (fixed) asymmetric costs. Overall, the evidence shows that asymmetry makes it much harder to collude in the lab; observed quantities are clearly more competitive than in symmetric configurations and often in line with static Nash predictions.

Most of this literature also looks at environments with complete information. Exceptions include Fouraker and Siegel (1963), Carlson (1967), and Mason and Phillips (1997), who have conditions in which a given player does not have any information about the payoff of the other player(s). Thus, in those studies, the games are not Bayesian, in the technical sense (Harsanyi, 1967) of having a prior distribution of types commonly known to all players.[5] Not all of those studies compare complete-information to incomplete-information environments; when they do, they report a tendency for collusion to be harder to achieve in the presence of incomplete information.

To our knowledge, no paper so far has looked at a (finitely) repeated standard Bayesian 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.

We conduct a series of Cournot laboratory experiments with two players under fixed matching and finite repetition. Our treatments include full symmetry and complete information (with two variants: constant costs over time and random realizations of symmetric costs every round), some slight 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 difference between the high and the low cost).

We uncover the following main findings. First, 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). This also applies to the treatment where firms have the same costs but those are drawn anew every period. Second, we find that as soon as any level of asymmetry or incomplete information about current-period costs is introduced, collusion disappears and observed average individual quantities are remarkably close to the static Bayes-Nash equilibrium values. We do not observe differences in collusion levels among Bayes-Nash 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 Offerman et al. (2002). Specifically, we introduce a conditional imitation process that consists for a given player in adopting the 'exemplary' choices made by the other player, i.e., choices which, if played in the relevant state, would increase the sum of (expected) payoffs for both players (when compared to the choices currently made by the player).[6] Simulations show that such an adjustment process converges towards collusive outcomes by contrast to standard best-response dynamics that lead to the Cournot (Bayes-Nash) outcomes.

We find evidence that in the treatments where either cost asymmetry or incomplete information about current cost conditions is present, subjects' adjustments are more in line with Cournot best-response to the opponent's previous choice rather than with imitation of 'exemplary' behavior by this opponent. By contrast, in the symmetric, complete-information treatments, players put less weight on playing a best response to their opponent's last (relevant) round choice and more weight on (conditionally) imitating it, and that allows them 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 treatments where there is either asymmetry

---

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

[5] For example, Mason and Phillips (1997) study behavior when players' payoffs were either common knowledge or private information. In the latter case, players were not given any information on the distribution of payoffs of their opponents.

[6] This process nests the 'follow-the-exemplary-other-firm' learning rule of Offerman et al. (2002) as a special case. For details see our section 4.2.

or incomplete information about current-period costs, 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 about current-period conditions.

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 so-called "coordinated effects" (i.e. the possibility of tacit collusion) in merger control. Davis et al. (2011) indeed show that the European Commission concerns itself with collusion threats only in the case of post-merger *symmetric* duopolies.

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 could 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 Bayesian Cournot environment brings a measure of stochasticity to a literature which has mainly focused on very stable (indeed, identically repeated) contexts.

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; in Section 3 we describe our experimental set-up and the various treatments; Section 4 contains our findings, the specification of learning dynamics and our analysis of subjects' adaptive behavior; and in Section 5 we discuss our results in the light of related literature.

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 equilibrium (BNE). It is routine to show that these quantities are given by:

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

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

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))

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

The parameters used in the experiment are provided in column 3 of Table 1 and the Bayes-Nash equilibrium predictions, given parameters, are provided in column 3 of Table 2.

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.[7]

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 necessarily agree on the collusive actions: Static joint profit maximization calls for the high-cost firm not to produce at all and for the low-cost firm to produce the monopoly quantity. If costs are drawn at random every period and the firms play those strategies for many periods, then they might, on average, collect half of the monopoly profit. However, a high-cost firm might be particularly wary that its rival might not reciprocate and cut production next time the cost configuration is reversed. If, for this reason, firms insist on producing the same quantity every period, then, because of the cost difference, they disagree about the optimal level.[8] Finally, in a proper Bayesian Cournot game where cost draws are privately observed, cooperation/collusion problems are arguably even more severe. In particular, it is no longer optimal, from the point of view of joint profits, to stop production when one draws a high-cost (as there is a non-zero probability that both firms draw such costs). We are not aware of a theoretical paper that solves for the optimal collusive scheme in this environment.[9]

---

[7] Note that in a repeated version of the Bayesian Cournot game, types are drawn anew in every period. Thus, there is no role for reputation building of the kind first shown by Kreps and Wilson (1982) in the case of fixed types, drawn once and for all at the beginning of the supergame.

[8] For more on this, see Schmalensee (1987) for a theoretical, and Fischer and Normann (2019) for a theoretical and experimental investigation of this case.

[9] For Bertrand games, the optimal collusive scheme has been characterized by Athey and Bagwell (2001) in the case of inelastic demand. For Cournot games, optimal schemes have been characterized only in the case of privately observed costs that are drawn once and for all, see Chakrabarti (2010). Under fixed types, if communication between firms is allowed, Roberts (1985) has pioneered a mechanism design approach to the problem.

Table 1

Experimental design.

Treatment	Info	Parameter Choices	#Subjects	#Markets	#Obs
		In all Treatments: $a = 120$ and $b = 1$			
30–C	C	$c_1 = c_2 = 30$	14	7	840
29-29-31-31–C	C	$\lambda_1 = \lambda_2 = 0.5$, $c_1 = c_2 = 29$, $c_1 = c_2 = 31$	36	18	2,160
29-31–C	C	$\lambda_1 = \lambda_2 = 0.5$, $c_1^H = c_2^H = 31$, $c_1^L = c_2^L = 29$	30	15	1,800
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
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
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
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 the column labeled "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 in the BNE."

Table 2

Summary statistics.

Treatment	Costs	Quantity in the BNE	Average Individual Quantity Observed		
			Rounds 1-30	Rounds 31-60	Rounds 1-60
30–C	$c = 30$	30	26.72**	25.72*	26.22**
			(1.26)	(1.24)	(1.09)
29-29–31-31–C	$c_1^L = 29$, $c_2^L = 29$	30.33	28.18**	28.24**	28.22**
			(0.76)	(0.77)	(0.72)
	$c_1^H = 31$, $c_2^H = 31$	29.67	27.44***	27.34**	27.41***
			(0.64)	(0.76)	(0.68)
29-31–C^a	$c_1^L = 29$, $c_2^L = 29$	30.33	28.81	30.27	29.69
			(0.79)	(0.71)	(0.66)
	$c_1^L = 29$, $c_2^H = 31$	31	30.75	30.68	30.70
			(0.78)	(0.47)	(0.58)
	$c_1^H = 31$, $c_2^L = 29$	29	28.15	29.19	28.72
			(0.80)	(0.54)	(0.64)
	$c_1^H = 31$, $c_2^H = 31$	29.67	28.49	29.75	28.99
			(0.64)	(0.50)	(0.56)
29-31–I	$c^L = 29$	30.5	29.72	29.26	29.49
			(0.85)	(0.95)	(0.88)
	$c^H = 31$	29.5	28.04	28.89	28.45
			(0.73)	(0.92)	(0.78)
25-35–I	$c^L = 25$	32.5	32.22	32.30	32.27
			(0.77)	(0.53)	(0.57)
	$c^H = 35$	27.5	27.63	28.25	27.93
			(0.41)	(0.70)	(0.49)
20-40–I	$c^L = 20$	35	34.52	35.69	35.12
			(0.71)	(0.77)	(0.65)
	$c^H = 40$	25	24.96	25.27	25.08
			(0.58)	(0.78)	(0.56)
20-40-10-50–I	$c_1^L = 20$	35	35.00	36.50	35.64
			(1.46)	(1.96)	(1.62)
	$c_1^H = 40$	25	24.18	23.44	23.77
			(1.45)	(1.17)	(1.24)
	$c_2^L = 10$	40	37.72	39.16	38.49
			(1.84)	(2.20)	(1.88)
	$c_2^H = 50$	20	19.86	17.90	18.93
			(1.04)	(1.17)	(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 29-31–C, BNE and observed quantities refer to those of player 1. Test statistics refer to two-tailed Wilcoxon 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 1%, 5%, 10% level, respectively.

3. Experimental design and procedures

In the experiment, subjects participated in 60 consecutive rounds of decision-making. 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 used 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, $c_i, i = 1, 2$ and with respect to the information about the cost structure in the market. In three out of the seven treatments (indicated with the letter 'C' in the treatment's name), the costs of both firms in a given round were common knowledge. In the other four treatments (with the letter 'I'), subjects only knew their own cost. More precisely, 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 assigned with probability 0.5. In 29-29–31-31–C, although the cost level is drawn randomly each period, the two firms receive the same costs. Symmetry is thus preserved. By contrast, in 29-31–C, firms' costs are *independently* drawn in each period and can therefore end up being different. While in 29-29–31-31–C and 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 all C–treatments and three of the four I–treatments (29-31–I, 25-35–I, 20-40–I) are *ex-ante* symmetric. Treatment 29-31–C could be *ex-post* asymmetric. Treatment 20-40–10-50–I is *ex-ante* and *ex-post* asymmetric as one firm has the two possible cost levels 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.

The comparison between 30–C and 29-31–C allows us to measure the net effect of introducing random costs, under complete information. However, two changes happen at the same time: first, current costs, i.e., costs in a given round, might be different; second, the future is now uncertain as future cost configurations are not known. Comparisons with 29-29–31-31–C allow us to disentangle the two effects, since in this treatment, current costs are always identical, while the future remains uncertain (although symmetric). The comparison between 29-31–C and 29-31–I allows us to single out the role of incomplete information, at unchanged cost structure. Comparisons among 29-31–I, 25-35–I, and 20-40–I allow us to detect any impact of the size of cost asymmetry, given incomplete information. Treatment 20-40–10-50–I allows us to detect any potential additional impact of *ex-ante* asymmetry.

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.[10] 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 costs (own cost in I-treatments or both costs in C-treatments), the quantity decisions of both firms, and their own profit in that 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 Web Appendix C 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 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 University Berlin and Humboldt University Berlin. Participants were students (33% female), mostly from economics, business, natural sciences, or engineering. Altogether, we conducted 95 markets with 190 subjects and collected 11,400 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. At the end of the experiment, subjects were paid on the basis of their cumulated earnings over the 60 rounds of play. The average earnings per subject in the experiment was 18.02 EUR.[11] Sessions took about 60 minutes to complete.

4. Experimental results

4.1. Aggregate results

Table 2 provides summary statistics for our experimental results. To account for statistical dependence of observations over time within a given market, 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.[12] Table 2 also shows the results of two-tailed Wilcoxon tests of whether the sample mean is equal to Bayes-Nash equilibrium values. The unit of observation for the tests is market averages of individual quantities. Looking at Table 2, we make a number of observations. First, for treatment 30–C, we find confirmation of the known result that subjects are, on average, partially able to collude.[13] For the three time intervals considered, the Wilcoxon test

---

[10] The profit calculator provides essentially the same information as commonly used payoff tables, but helps to avoid a possible bias due to limited cognitive abilities of participants (Huck et al. 2000, p. 42). Due to its availability in all treatments, the use of the payoff calculator cannot explain the treatment effects we report in this paper. Note that Requate and Waichman (2011) report that "the most standard variations, which are the use of a profit table or a profit calculator, yield indistinguishable performance." (p. 36).

[11] In addition to their earnings in the experiment, subjects were given an initial (show-up) payment of 2.50 EUR in all treatments except 29-29–31-31–C, and 6 EUR in treatment 29-29–31-31–C (due to a change in the lab rules at the time the latter treatment was conducted).

[12] As mentioned before, perfectly collusive outcomes are not well-defined in all treatments but 30–C due to the asymmetry of interests. Hence, we do not provide collusion indices (Friedman, 1971) as is customary in many papers on market experiments.

[13] Note that the individual perfectly collusive quantity (half of the monopoly quantity) is 22.5.

indicates that the observed individual market averages are statistically significantly below the Nash equilibrium. Second, in treatment 29-29–31-31–C, subjects also collude on average, though to a lesser extent than subjects in treatment 30–C. Again, Wilcoxon tests for treatment 29-29–31-31–C indicate that the observed individual market averages are statistically significantly below Bayes-Nash equilibrium levels for all time horizons considered in Table 2. Third, in all other treatments the observed averages are remarkably close to the Bayes-Nash equilibrium, and in none of the cases does the Wilcoxon test reject equality of observed averages with predicted values ($p > 0.1$).[14] This is perhaps most surprising in treatments 29-31–C and 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 A.1 in Web Appendix B.1 shows the distributions (histograms) of averages of individual quantities per market for each treatment separately.

The question is why we observe successful collusion in the case of full symmetry and complete information but neither in complete-information treatment 29-31–C nor in any of the incomplete-information treatments (one of which involves only minute payoff differences). The results for 29-29–31-31–C, where subjects manage to achieve some collusion, while they fail to do so in 29-31–C, suggest that (small, symmetric) uncertainty about the future is not the main driver of the breakdown in cooperation. Asymmetry or uncertainty about *current* cost conditions, the fact that subjects do not know that they have identical interests in the current period, seems to play a key role. We thus conjecture that in treatments in which asymmetry or incomplete information about current cost conditions is present, players are less willing to cooperate. This could translate into some unwillingness to go along with opponents who try and cut output.

In the next subsection, we shed light on this issue by analyzing to what extent behavior in our treatments accords with several learning dynamics.

4.2. Learning dynamics: theory

To understand subjects' quantity choices, we attempt to compare their sequence of decisions to some learning rules that could theoretically explain how their behaviors evolve over time. Two caveats apply. First, many learning rules or models have been proposed in the game-theoretic and experimental literatures over the past 40 years so that any focus on a subset of them can always be construed as arbitrary. One is often led to use those rules that have been extensively studied or have interesting theoretical properties or have been shown to have some predictive power in a number of contexts. We will be no exception.

Second, most of the learning-in-games literature has been concerned with stable and symmetric environments where players have the same strategy set and the same payoff function in every period. When costs are independently and randomly drawn, our Bayes-Nash Cournot stage game is *ex-ante* symmetric (except in 20-40–10-50–I, which we leave aside), but in any given period, because of the cost draws, two players in a market may happen to have different costs and therefore different payoffs. Moreover, the situation might be different in the next period. This considerably complicates the specification of "reasonable" learning rules. To illustrate, suppose that in treatment 29-31–C, in period t, both player 1 and player 2 drew c_L. Suppose further that in period $t + 1$, player 1 drew cost c_L again while player 2 drew c_H. Imagine you want to specify a simple best-response rule where a player plays her best response to the latest relevant quantity played by her opponent. In $t + 1$, should player 1 best-respond to the quantity chosen by player 2 in period t? That is defensible but not obvious. Indeed, player 1 knows that player 2 now has a different payoff and may be led to play another quantity than the one she last chose. She may want to react to the quantity chosen by player 2 the last time she drew a cost of c_H, which may not be the previous period in general and is not the previous period in this specific example.

We could not find relevant guidance in the literature about such questions and about learning in changing environments, more generally. Hence, our study of individual behavior is a first pass at dealing with the issue. We chose to adapt some existing rules to our context (and later check that using some other, perhaps more "naive", rules also supports our analysis) but we do not claim any of those to be of universal application. We believe our choices are sensible but the general question of learning in dynamic games remains open.

We focus on two learning rules: best-response to relevant past play and conditional imitation of the other player's actions. Those two processes have the interesting property of converging toward Cournot outcomes and collusive outcomes (in some specific sense which we explain below), respectively.

4.2.1. Specification of the best-response ("BR") process

In some of our treatments, in each period, players face a potentially different (observable) cost configuration. We define the first learning rule as best-responding to the last "relevant" quantity played by the other player. The general idea is as follows. If $r_i(s)$, where s is a particular cost configuration, stands for the latest quantity played by firm j in that cost configuration, firm i will best-respond to $r_i(s)$ next time cost configuration s is drawn. It is as if player i were keeping mental track of the quantities last chosen by player j in the various cost configurations.

For example, in treatment 29-31–C, player i observes the cost type of j on top of her own cost. Thus, there are four (common) cost configurations or states, and the state space is $S = \{(c_L, c_L), (c_L, c_H), (c_H, c_L), (c_H, c_H)\}$, where the first variable stands for the cost drawn by player 1 and the second, for the cost drawn by player 2. Let $b(\cdot)$ denote the Cournot best-response operator. Let q_j^t be the quantity played by player j in period t. In any period, an updating process takes place. If player i finds herself in state $s \in S$ in

---

[14] The existing literature on Cournot markets with asymmetric costs, complete information and fixed matching reports observed average individual quantities to be close to static Nash predictions (see Mason et al., 1992; Fonseca et al., 2005; or Normann et al., 2014).

period t, she will play $b(r_i(s))$ and then update $r_i(s)$ to correspond to q_j^t so that next time cost configuration s arises again, she will best-respond to q_j^t.

In incomplete-information treatments, for example, 29-31–I, player i observes only her own cost level. Hence, the set of observable cost configurations is $S = \{c_L, c_H\}$. That is, when i observes a particular cost level for herself, she will best-respond to the quantity chosen by j last time she happened to have drawn that cost level. That is the sense in which a player best-responds to the quantity chosen by the other player in the last "relevant" round.

4.2.2. Specification of the conditional imitation ("CI") process

Imitation has been shown to play a role in a number of experimental games (see, e.g., Huck et al., 1999, 2002, Rassenti et al., 2000; Offerman et al., 2002). In an environment where players do not necessarily have the same payoff, it is however not obvious how to specify an imitation process (or even, whether imitation should take place in the first instance). We specify an imitation process that is not "naive" (reproducing the choices of the other player in all circumstances) but is concerned with the "exemplary" choices made by the other player.[15] Exemplarity is defined with reference to the maximization of joint payoffs: the quantity chosen by the other player in a particular cost configuration is exemplary if, when a player found herself in a comparable situation, she would find it better, from the point of view of the expected sum of profits, to play the quantity chosen by the other player rather than the quantity that she has last played in that state.

Thus, each player keeps track of a set of "exemplary" quantities, one for each cost configuration at which she may be called upon making a decision. At the end of every period, given the current cost configuration and play, a player asks herself whether the quantity just chosen by the other player would be a good idea for her to play in comparable circumstances. If (and only if) the answer is yes, then that player will update her list of exemplary quantities and imitate the other player next time those relevant circumstances arise. That is the sense in which imitation is "conditional".

This process is inspired by the "follow-the-exemplary-other-firm" process put forward by Offerman et al. (2002) in the case of symmetry, identical repetition, and complete information, and actually nests that case as a special (degenerate) case with one (common) cost configuration.

To be more specific, in the asymmetric case under complete information (treatment 29-31–C), player i observes the cost type of j on top of her own cost. Thus, there are four (common) states: $(c_L, c_L), (c_L, c_H), (c_H, c_L), (c_H, c_H)$. So, each player i keeps track of four exemplary quantities: $\{\mathbf{q}_i(c_L, c_L), \mathbf{q}_i(c_L, c_H), \mathbf{q}_i(c_H, c_L), \mathbf{q}_i(c_H, c_H)\}$.

In the symmetric states $\{(c_L, c_L), (c_H, c_H)\}$, at the end of the period, player i asks herself whether it would have been a good idea for her to play the quantity just chosen by j. If so, she updates the relevant state variable and will play that quantity next time the cost configuration arises again.

In the asymmetric states $\{(c_L, c_H), (c_H, c_L)\}$, at the end of the period, player i asks herself whether it would have been a good idea to play the quantity just chosen by j, had she been in j's position, that is, had the two roles (cost types) been switched. That is consistent with the fact that player i understands that she has just had a different cost draw than player j and that player j's choice is relevant to her, from the point of view of joint profit maximization, only in the reversed cost configuration. If the answer is yes, she will update the relevant state variable and play that "exemplary" quantity next time roles are switched.

In the case of incomplete information, player i is not aware of player j's cost and so this issue does not arise, as i simply asks herself whether the quantity chosen by j would be a good idea, from the point of view of joint profit maximization, for a player who has just drawn her cost level and does not know the cost of her opponent for sure. We provide details about the updating processes in Web Appendix A.[16]

Note that under this conditional imitation rule, players move towards quantities that are more and more "collusive" since they choose to change their behavior ("update") only when this is better from the point of view of joint expected profit-maximization. By contrast, as is intuitive, best-response dynamics in Cournot environments lead players towards Cournot equilibrium outcomes.

4.2.3. Simulation results

That is corroborated by our simulations of (stochastic versions of) best-response and conditional imitation dynamics for all treatments (see Table A.1 in Web Appendix A.3). First, best-response dynamics converge to the Bayes-Nash equilibrium of the stage game in all treatments.[17] Second, for the complete-information treatments, conditional imitation dynamics converge to the solution to joint profit maximization, which consists of having the low-cost firm produce the monopoly quantity (for that cost level), while the high-cost firm produces nothing. In incomplete-information treatments, play also converges to quantities that are below the Bayes-Nash predictions and decreasing in costs, but both strictly positive. Note that in those latter treatments, the sum of the two limit quantities is roughly equal to 45, which is the monopoly output in Treatment 30–C.

Hence, it appears that players following best-response dynamics ("BR") would converge towards "competitive" play (Bayes-Nash Cournot outcomes) while players following conditional imitation dynamics ("CI") would achieve collusive outcomes. Thus,

---

[15] Note that a deterministic unconditional imitation process where players simply play the quantity chosen by the other player in the latest round would not necessarily converge, as players would take turns in playing the two initially-chosen quantities.

[16] We leave aside treatment 20-40–10-50–I. In this treatment, players know for a fact that they will never have the same cost level as their opponent's. Thus, the very idea of imitating the other player's behavior is called into question.

[17] For complete-information Cournot duopoly with linear demand and costs, this is of course known since Theocharis (1960). With costs randomly drawn every period, the quantity played varies from one period to the next. So, q_i^t, as a series, does not technically converge. Quantities played as a function of the cost configuration converge.

Table 3
Summary of hypothesis tests for adjustment dynamics per individual.

	H_0	H_1	30-30–C	29-29–31-31–C	29-31–C	29-31–I	25-35–I	20-40–I
			Percentage of subjects for which H_0 is rejected at the 5% level (in favor of H_1)					
"Previous relevant	$p = 0.5$	$p < 0.5$	57.14	33.33	6.67	21.43	7.14	7.69
rounds"	$p = 0.5$	$p > 0.5$	28.57	47.22	60.00	60.71	71.43	53.85

Notes: This table shows the results of binomial tests at the individual level, using all data. Note that $p < 0.5$ ($p > 0.5$) means that behavior is closer to conditional imitation (best response).

determining whether players are closer to using BR learning rules, than to using CI, potentially allows us to explain why outcomes may be more competitive in some treatments than in some others.

4.3. Learning dynamics: results

In order to explore the learning patterns in our experiment, we compare the observed adjustment behavior with those predicted by the two learning dynamics discussed in the previous section. Fig. 1 shows the evolution of the average observed differences $\Delta q^{OBS} = q_i^t - q_i^{t-1}$ (solid line), where q_i^t and q_i^{t-1} are player i's quantity choices in current period t and previous relevant period $t-1$, respectively, averaged across cost configurations and markets per treatment. The two other lines in Fig. 1 indicate the predicted average differences $\Delta q^{BR} = BR_i^{t-1} - q_i^{t-1}$ (dotted line) and $\Delta q^{CI} = CI_i^{t-1} - q_i^{t-1}$ (dashed line), where BR_i^{t-1} and CI_i^{t-1} are the point predictions implied by playing "best response" and "conditional imitation," respectively.[18] Inspecting Fig. 1, we make one main observation: In treatments 30–C and 29-29–31-31–C, the line representing Δq^{OBS} is between the lines representing Δq^{BR} and Δq^{CI}. However, in all other treatments, the lines representing Δq^{OBS} and Δq^{BR} are very close to each other, while the line representing Δq^{CI} is clearly and substantially below the two other lines. This suggests that in treatment 29-31–C and all incomplete-information treatments, behavior is much more in line with best-response adaptations than with conditional-imitation behavior.

We performed non-parametric tests to verify that behavior in the treatments with symmetric costs and complete information (as shown in the two top panels in Fig. 1) is qualitatively clearly different than behavior in the asymmetric-cost treatment 29-31–C and the incomplete-information treatments (as shown in the other panels in Fig. 1). For this purpose, we computed averages across all periods at the individual market level of the terms Δq^{OBS}, Δq^{BR} and Δq^{CI}. Table A2 in Web Appendix B.1 shows market averages (with standard errors of the mean in parentheses) of these terms for each possible cost configuration per treatment. Table A2 also shows the results of two-tailed Wilcoxon tests of whether the sample mean of Δq^{OBS} is the same as either Δq^{BR} or Δq^{CI}, respectively. The unit of observation for the tests is market averages for each treatment. For treatments 30–C and 29-29–31-31–C, we find that the market averages of Δq^{BR} are significantly larger than the market averages of Δq^{OBS} and the latter significantly larger than the market averages of Δq^{CI}. For all other treatments, we find that the market averages of Δq^{OBS} and Δq^{CI} are still significantly different, but that the market averages of Δq^{OBS} and Δq^{BR} are statistically indistinguishable from each other. This indicates that adaptation behavior in treatments other than 30–C and 29-29–31-31–C are in line with best-response behavior and clearly different from conditional-imitation behavior. Average observed adaptation behavior in the treatments 30–C and 29-29–31-31–C appears to be a mix between BR and CI. Indeed, in those treatments we find that subjects in some markets successfully collude, while others rather play according to Nash equilibrium predictions.

We also conducted tests at the individual level. To do so, we define the following dummy variable labeled "INDEX" for each individual decision observed in the experiment:
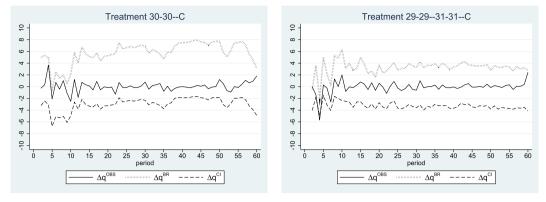
$$INDEX = \begin{cases} 1 & \text{if } |\Delta q^{OBS} - \Delta q^{BR}| < |\Delta q^{OBS} - \Delta q^{CI}| \\ 0 & \text{otherwise.} \end{cases} \tag{4}$$
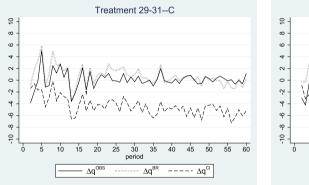
That is, variable INDEX is equal to 1 if the observed adaptation (Δq^{OBS}) is closer to the one prescribed by best-response behavior (Δq^{BR}) rather than conditional-imitation behavior (Δq^{CI}), and 0 otherwise. Under the assumption that the BR and CI dynamics explain observed behavior equally well, the variable INDEX as defined in (4) should be binomially distributed with $p = 0.5$ for each subject.[19] Note that an observed $p > 0.5$ ($p < 0.5$) means that a subject's behavior is closer to BR (CI). The results are presented in Table 3, where we ignore treatment 20-40–10-50–I, see footnote 16. The entries in this table indicate the share of subjects per treatment for which H_0: $p = 0.5$ is rejected. We make the following main observations. In treatment 30–C, the percentage of subjects for which H_0 is rejected in favor of H_1: $p < 0.5$ is clearly *larger* than the percentage of subjects for which H_0 is rejected in favor of H_1: $p > 0.5$. In the other treatments, the percentage of subjects for which H_0 is rejected in favor of H_1: $p < 0.5$ is clearly *smaller* than the percentage of subjects for which H_0 is rejected in favor of H_1: $p > 0.5$. This indicates that at the individual level, in all treatments but 30–C subjects' adjustments are on average more in line with best-response behavior than with conditional-imitation behavior, highlighting the effect of introducing asymmetric costs and incomplete information.
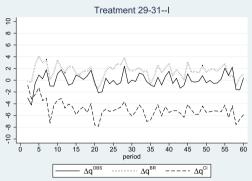
We also investigate the above results by means of extensive regression analysis. The variable INDEX as defined in (4) serves as the dependent variable in probit panel regressions involving the following independent variables: A binary variable ("AsymCosts")
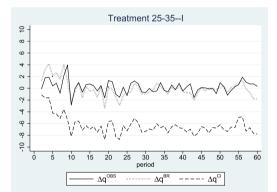
---

[18] For the initialization of the CI dynamics, we use the average across markets of quantities chosen by the other player in the relevant cost configuration in period 1.
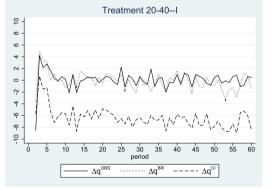
[19] We hasten to acknowledge that this approach (counterfactually) relies on the observations in a given market being independent from one period to the next.
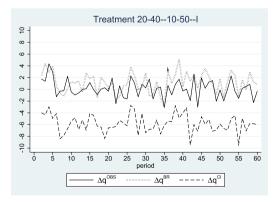
Notes: The panels in this figure show the evolution of the terms $\Delta q^{\text{OBS}} = q_i^t - q_i^{t-1}$ (solid line), $\Delta q^{\text{BR}} = BR_i^{t-1} - q_i^{t-1}$ (dotted line) and $\Delta q^{\text{CI}} = CI_i^{t-1} - q_i^{t-1}$ (dashed line) per treatment (see the definitions in the text, starting on page 8), averaged across cost configurations and markets of the same treatment.

Fig. 1. Evolution of average observed and predicted changes in quantities.

indicating whether cost asymmetry exists, a binary variable ("PrivInfo") indicating the presence of incomplete information, and a binary variable ("RandSym") indicating randomness in cost assignment in a symmetric and complete-information treatment. The regression results, reported in Table A3 in Web Appendix B.2, show that "AsymCosts" and "PrivInfo" are positive and statistically significant, while "RandSym" is also positive but insignificant. Given the definition of the dependent variable in (4), these results confirm that the presence of cost asymmetry or incomplete information significantly tilts behavior towards best-responding. In the Web Appendix B.2, we show that other definitions of cost asymmetry (ex post asymmetry) or conditioning on the last round played rather than the last "relevant" round leave this result unchanged. We also perform robustness checks by (i) changing the initial conditions of the CI process (predicted BNE quantities instead of the average observed quantities played in period 1), (ii) redefining the CI process by allowing players to learn not only from the quantities chosen by the other player but also from the ones they have just chosen in the current round, and (iii) replacing the CI process with unconditional imitation (UI) of the quantity played by the other player in the last relevant round. Although the level of significance of our regressors occasionally changes from one regression to the next, all specifications indicate that cost asymmetry and complete information foster best-response behavior (see Web Appendix B.3).

Finally, the analysis of the recorded simulations conducted by subjects with the help of the payoff calculator prior to the actual quantity choices confirms the observed difference in the subjects' decision approach used in treatment 30–C compared to the treatments with either cost asymmetry or incomplete information.[20] See Table A5 in Web Appendix B.4 for detailed results. Specifically, it is apparent that subjects in treatment 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 30–C.

5. Concluding remarks

5.1. Summary

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 about current costs (which are either constant over time or drawn anew every period), (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 symmetry of and complete information about current cost conditions, 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, collusion breaks down and observed average individual quantities get remarkably close to the static (Bayesian) Nash equilibrium values. This is so despite a high number of repetitions (60 rounds) and fixed matching.

The results of the analysis of 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 conditional imitation in those treatments involving asymmetry or incomplete information about current cost conditions. This explains our results since simulations show that best-response dynamics converge to static (Bayesian) Nash equilibrium quantities, whereas conditional-imitation dynamics converge to collusive outcomes.

The standard adjustment dynamics that are present in the literature have been developed in the context of symmetric and complete-information games. Adapting them to asymmetric or uncertain environments is not necessarily obvious and we have made a first pass at it. Other learning rules and other approaches can be conceived and we do not claim that the ones we used in this paper are the best predictors of subjects' behavior. We think that future theoretical and empirical work should probe whether alternative specifications can better account for players' adaptations over time in such environments.

5.2. Relation to the literature

In their Bertrand duopoly treatment (as well as the ones with 3 or 4 firms) with costs independently and repeatedly drawn from a common distribution, 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 et al., 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.

When it comes to collusion, the Cournot model can be interpreted as an extended form of a Prisoner's Dilemma (PD). There are several papers that explore the role of stochasticity, asymmetry and incomplete information in (finitely) repeated PD games. See Andreoni and Miller (1993), Bereby-Meyer and Roth (2006), Ahn et al. (2007), or Zhang et al. (2022), to name a few.[21] However,

---

[20] 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 (own and of the opponent) and were then shown the resulting profit for themselves.

[21] For a survey of the literature on (in)finitely repeated PD games, see Embrey et al. (2018), Bo and Fréchette (2018), Mengel (2018).

these papers, like the Cournot literature, implement incomplete information in a context that is quite different from the one of Bayesian games (Harsanyi, 1967). For example, Andreoni and Miller (1993) manipulate subjects' beliefs about the opponent's type by varying the probability of interacting with a computerized opponent. Bereby-Meyer and Roth (2006) compare behavior in PD games with either deterministic or 'noisy' (as opposed to type-contingent) payoffs. Notwithstanding those differences, similarly to our findings, these studies tend to show that cooperation is harder to sustain when a (potential) difference in players' payoffs leads to an increase in the uncertainty about the opponents' intentions.

There is also a sizeable experimental literature studying collusive behavior in auctions, which stand for prime examples of Bayesian games. However, by construction, this literature does not compare complete to incomplete information.[22] Moreover, very few studies concern themselves with asymmetry and, when doing so, typically focus on *ex-ante* asymmetry in the distribution of valuations (see, e.g., Güth et al., 2005).[23] Instead, researchers compare various auction formats or various 'institutions' thought to affect collusion.[24] The most closely related studies in this strand of literature are the ones of Sherstyuk (1999) and Sherstyuk (2002) which investigate whether and to what extent varying the degree of bidders' value (*ex-post*) asymmetry and the gains from collusion in oral ascending auctions affects collusion. Note that in ascending auctions, players' actions are observable and reaction to opponents' actions are possible before the game ends. In that sense, Bayes-Cournot games are comparable to sealed-bid auctions rather than ascending auctions but repetition allows for some reactions over time. Those differences notwithstanding, Sherstyuk (1999) demonstrates that if all bidders have the same value for the object (the special case of symmetry and complete information), then collusive outcomes are sustainable only in oral ascending auctions, whereas outcomes in sealed-bid auctions are significantly more competitive. Sherstyuk (2002) shows that, if bidders' values are private information, but drawn from the same (*ex-ante* symmetric) distribution, increasing the level of bidders' value *ex-post* asymmetry leads to an increase in market competitiveness. Those results are broadly in line with our findings.

Finally, our results are also reminiscent of those in Crawford et al. (2008). These authors report that in games with symmetric payoffs salient labels generate high coordination rates, while the effectiveness of salient labels is significantly reduced in the presence of even slight payoff differences between players. Crawford et al. (2008) mainly invoke level-k thinking to explain their one-shot experiments.[25]

It is worth noting that our experiment was concerned with decision-making in isolation, without the possibility for subjects to communicate. Communication is often reported to help sustain cooperation in social dilemmas. For instance, in the context of symmetric Bertrand oligopolies, Fonseca and Normann (2012) show that pre-play communication increases profits for any number of firms. Fischer and Normann (2019) show that in asymmetric Cournot duopolies, talking helps reduce output. Agranov and Yariv (2018) show that communication reliably facilitates collusion in one-shot sealed-bid auctions. This begs the question as to whether communication would also restore cooperation in our Cournot-Bayes-Nash environments. We plan to investigate this matter in future work.

Declaration of competing interest

We declare that there are no financial and personal relationships with other people or organizations that could inappropriately influence (bias) our work.

Data availability

Data will be made available on request.

Appendix A. Supplementary material

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.geb.2023.12.004.

References

Abbink, K., Brandts, J., 2005. Price competition under cost uncertainty: a laboratory analysis. Econ. Inq. 43, 636–648.
Agranov, M., Yariv, L., 2016. Collusion through communication in auctions. Games Econ. Behav. 107 (C), 93–108.
Ahn, T.K., Lee, M., Ruttan, L., Walker, J., 2007. Asymmetric payoffs in simultaneous and sequential Prisoner's Dilemma games. Public Choice 132, 353–366.
Altavilla, C., Luini, L., Sbriglia, P., 2006. Social learning in market games. J. Econ. Behav. Organ. 61, 632–652.

---

[22] Complete information in an independent private value auction would reduce to standard Bertrand competition.

[23] Exceptions include Avery and Kagel (1997) who compare symmetric and asymmetric payoffs in second-price common-value auctions and Andreoni et al. (2007) who study asymmetric information about opponents' types in standard (ex-post) private-value auctions.

[24] For example, Hu et al. (2011) compare behavior in single-unit auctions across three different auction formats allowing for explicit collusion and introducing *ex-ante* asymmetry among groups of bidders. Kwasnica and Sherstyuk (2007) show that in repeated multi-object ascending auctions, collusion can be achieved through bidders' coordination on payoff-superior outcomes. Hinloopen and Onderstal (2014) study the role of external enforcement and cartel detection, Agranov and Yariv (2016), Noussair and Seres (2020) focus on the role of communication.

[25] Interestingly, some of their results can only be explained by "team reasoning" where "players begin by asking themselves, independently, if there is a decision rule that would be better for both than individualistic rules, if both players followed the better rule" (p. 1448). Note that this kind of team reasoning is also the basis of the conditional-imitation approach suggested by Offerman et al. (2002) and used in our analysis.

Andreoni, J., Miller, J.H., 1993. Rational cooperation in the finitely repeated Prisoner's Dilemma: experimental evidence. Econ. J. 103, 570–585.

Andreoni, J., Che, Y.-K., Kim, J., 2007. Asymmetric information about rivals' types in standard auctions: an experiment. Games Econ. Behav. 59 (2), 240–259.

Avery, C., Kagel, J.H., 1997. Second-price auctions with asymmetric payoffs: an experimental investigation. J. Econ. Manag. Strategy 6 (3), 573–603.

Athey, S., Bagwell, K., 2001. Optimal collusion with private information. Rand J. Econ. 32 (3), 428–465.

Bereby-Meyer, Y., Roth, A.E., 2006. The speed of learning in noisy games: partial reinforcement and the sustainability of cooperation. Am. Econ. Rev. 96, 1029–1042.

Bigoni, M., Potters, J., Spagnolo, G., 2019. Frequency of interaction, communication and collusion: an experiment. Econ. Theory 68, 827–844.

Binger, B.R., Hoffman, E., Libecap, G.D., Shachat, K., 1990. An Experimetric Study of the Cournot Theory of Firm Behavior. University of Arizona. Working paper.

Bo, P.D., Fréchette, G.R., 2018. On the determinants of cooperation in infinitely repeated games: a survey. J. Econ. Lit. 56 (1), 60–114.

Carlson, J.A., 1967. The stability of an experimental market with a supply-response lag. South. Econ. J. 33, 305–321.

Chakrabarti, S.K., 2010. Equilibrium in Cournot oligopolies with unknown costs. Int. Econ. Rev. 51 (4), 1209–1238.

Crawford, V.P., Gneezy, U., Rottenstreich, Y., 2008. The power of focal points is limited: even minute payoff asymmetry may yield large coordination failures. Am. Econ. Rev. 98 (4), 1443–1458.

Davis, D.D., 2002. Strategic interactions, market information and predicting the effects of mergers in differentiated product markets. Int. J. Ind. Organ. 20, 1277–1312.

Davis, D.D., Reilly, R.J., Wilson, B.J., 2003. Cost structures and Nash play in repeated Cournot games. Exp. Econ. 6, 209–226.

Davies, S., Olczak, M., Coles, H., 2011. Tacit collusion, firm asymmetries and numbers: evidence from EC merger cases. Int. J. Ind. Organ. 29, 221–231.

Embrey, M., Fréchette, G.R., Yuksel, S., 2018. Cooperation in the finitely repeated Prisoner's Dilemma. Q. J. Econ. 133 (1), 509–551.

Engel, C., 2007. How much collusion? A meta-analysis of oligopoly experiments. J. Compet. Law Econ. 3 (4), 491–549.

Feinberg, R.M., Husted, T.A., 1993. An experimental test of discount-rate effects on collusive behaviour in duopoly markets. J. Ind. Econ. 41, 153–160.

Fischbacher, U., 2007. z-Tree: Zurich toolbox for ready-made economic experiments. Exp. Econ. 10, 171–178.

Fischer, C., Normann, H.-T., 2019. Collusion and bargaining in asymmetric Cournot duopoly–an experiment. Eur. Econ. Rev. 111, 360–379.

Fonseca, M.A., Normann, H.-T., 2012. Explicit vs. tacit collusion—the impact of communication in oligopoly experiments. Eur. Econ. Rev. 56 (8), 1759–1772.

Fonseca, M.A., Huck, S., Normann, H.-T., 2005. Playing Cournot although they shouldn't: endogenous timing in experimental duopolies with asymmetric cost. Econ. Theory 25, 669–677.

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

Friedman, J.W., 1971. A non-cooperative equilibrium for supergames. Rev. Econ. Stud. 38, 1–12.

Güth, W., Ivanova-Stenzel, R., Wolfstetter, E., 2005. Bidding behavior in asymmetric auctions: an experimental study. Eur. Econ. Rev. 49, 1891–1913.

Harsanyi, J.C., 1967. Games with incomplete information played by "Bayesian" players, I—III. Part I. The basic model. Manag. Sci. 14, 159–182.

Hinloopen, J., Onderstal, S., 2014. Going once, going twice, reported! Cartel activity and the effectiveness of antitrust policies in experimental auctions. Eur. Econ. Rev. 70, 317–336.

Hoggatt, A.C., 1959. An experimental business game. Behav. Sci. 4, 192–203.

Hu, A., Offerman, T., Onderstal, S., 2011. Fighting collusion in auctions: an experimental investigation. Int. J. Ind. Organ. 29, 84–96.

Huck, S., Normann, H.-T., Oechssler, J., 1999. Learning in Cournot oligopoly-an experiment. Econ. J. 109, C80–C95.

Huck, S., Normann, H.-T., Oechssler, J., 2000. Does information about competitors' actions increase or decrease competition in experimental oligopoly markets? Int. J. Ind. Organ. 18, 39–57.

Huck, S., Normann, H.-T., Oechssler, J., 2002. Stability of the Cournot process – experimental evidence. Int. J. Game Theory 31, 123–136.

Huck, S., Normann, H.-T., Oechssler, J., 2004. Two are few and four are many: number effects in experimental oligopolies. J. Econ. Behav. Organ. 53, 435–446.

Kwasnica, A.M., Sherstyuk, K., 2007. Collusion and equilibrium selection in auctions. Econ. J. 117 (516), 120–145.

Kreps, D.M., Wilson, R., 1982. Reputation and imperfect information. J. Econ. Theory 27 (2), 253–279.

Mason, C.F., Phillips, O.R., Nowell, C., 1992. Duopoly behavior in asymmetric markets: an experimental evaluation. Rev. Econ. Stat. 74 (4), 662–670.

Mason, C.F., Phillips, O.R., Redington, D.B., 1991. The role of gender in a non-cooperative game. J. Econ. Behav. Organ. 15, 215–235.

Mason, C.F., Phillips, O.R., 1997. Information and cost asymmetry in experimental duopoly markets. Rev. Econ. Stat. 79 (2), 290–299.

Mengel, F., 2018. Risk and temptation: a meta-study on Prisoner's Dilemma games. Econ. J. 128, 3182–3209.

Mermer, A.G., Müller, W., Suetens, S., 2021. Cooperation in infinitely repeated games of strategic complements and substitutes. J. Econ. Behav. Organ. 188, 1191–1205.

Normann, H.-T., Requate, T., Waichman, I., 2014. Do short-term laboratory experiments provide valid descriptions of long-term economic interactions? A study of Cournot markets. Exp. Econ. 17, 371–390.

Noussair, C., Seres, G., 2020. The effect of collusion on efficiency in experimental auctions. Games Econ. Behav. 119, 267–287.

Potters, J., Suetens, S., 2009. Cooperation in experimental games of strategic complements and substitutes. Rev. Econ. Stud. 76 (3), 1125–1147.

Offerman, T., Potters, J., Sonnemans, J., 2002. Imitation and belief learning in an oligopoly experiment. Rev. Econ. Stud. 69, 973–997.

Raab, P., Schipper, B.C., 2009. Cournot competition between teams: an experimental study. J. Econ. Behav. Organ. 72, 691–702.

Rassenti, S., Reynolds, S.S., Smith, V.L., Szidarovszky, F., 2000. Adaptation and convergence of behavior in repeated experimental Cournot games. J. Econ. Behav. Organ. 41, 117–146.

Requate, T., Waichman, I., 2011. "A profit table or a profit calculator?" A note on the design of Cournot oligopoly experiments. Exp. Econ. 14, 36–46.

Roberts, K., 1985. Cartel behaviour and adverse selection. J. Ind. Econ. 33 (4), 401–413.

Sauerman, H., Selten, R., 1959. Ein oligopolexperiment. Z. Gesamte Staatswiss. 115, 427–471.

Schmalensee, R., 1987. Competitive advantage and collusive equilibria. Int. J. Ind. Organ. 5, 351–367.

Selten, R., Mitzekewitz, M., Uhlich, G.R., 1997. Duopoly strategies programmed by experienced players. Econometrica 65, 517–555.

Sherstyuk, K., 1999. Collusion without conspiracy: an experimental study of one-sided auctions. Exp. Econ. 2, 59–75.

Sherstyuk, K., 2002. Collusion in private value ascending price auctions. J. Econ. Behav. Organ. 48, 177–195.

Suetens, S., Potters, J., 2007. Bertrand colludes more than Cournot. Exp. Econ. 10 (1), 71–77.

Theocharis, R.D., 1960. On the stability of the Cournot solution on the oligopoly problem. Rev. Econ. Stud. 27, 133–134.

Waichman, I., Requate, T., Siang, C.N.K., 2014. Communication in Cournot competition: an experimental study. J. Econ. Psychol. 42, 1–16.

Zhang, L., Jin, Y., Xia, L., Xu, B., Abdullah, S.M.S., 2022. The effects of social distance and asymmetric reward and punishment on individual cooperative behavior in dilemma situations. Front. Psychol. https://doi.org/10.3389/fpsyg.2022.816168.