Discussion Paper No. 1200

ISSN (Print) 0473-453X
ISSN (Online) 2435-0982

AN INDIVIDUAL EVOLUTIONARY LEARNING MODEL MEETS COURNOT

Jasmina Arifovic
Liang Diao
Nobuyuki Hanaki

January 2023

The Institute of Social and Economic Research
Osaka University
6-1 Mihogaoka, Ibaraki, Osaka 567-0047, Japan

An individual evolutionary learning model meets Cournot*

Jasmina Arifovic† Liang Diao‡ Nobuyuki Hanaki§

January 3, 2023

Abstract

In this paper, we extend the individual evolutionary learning model by incorporating other-regarding considerations and apply the model to some Cournot games. Using the model fitted to the experimental data of a repeated 3-player Cournot game (with nonlinear cost and demand functions), we construct out-of-sample predictions regarding the "feedback effects" and "number effects" and test these using data gathered via newly conducted experiments. The prediction regarding the feedback effect is only partially confirmed, being observed for 3- and 4-player games but not the 2-player game. The prediction regarding the number effect is also partially confirmed in that while the model predicts the number effect to be observed with detailed and not aggregate feedback, the effect is observed with both types of feedback.

Keywords: Individual Evolutionary Learning, Oligopoly, Experiment

JEL Classification: D43, D83, D90

---

*We sincerely acknowledge the contribution of the late Jasmina Arifovic, who sadly passed away before her final editing of this then unfinished paper. Comments and suggestions from Nick Vriend are gratefully acknowledged. Part of this research was conducted when Hanaki was affiliated with the Université Côte d'Azur. This research has benefited from the financial support of (a) the Joint Usage/Research Center, the Institute of Social and Economic Research, Osaka University, (b) Grant-in-aid for Scientific Research Nos. 18K19954 and 20H05631 from the Japan Society for the Promotion of Science, and (c) the French National Research Agency (ANR) under an ORA-Plus project "BEAM" (ANR-15-ORAR-0004), and was performed within the framework of Programme Investissements d'Avenir under and UCA^{JEDI} (ANR-15-IDEX-01) of Université Côte d'Azur (project UCAinACTION).

†Simon Fraser University, Canada.

‡Simon Fraser University, Canada. E-mail: liang_diao@sfu.ca

§Institute of Social and Economic Research, Osaka University, Japan. E-mail: nobuyuki.hanaki@iser.osaka-u.ac.jp

1 Introduction

Cournot games have long been one of the workhorse models of industrial organization. One of their interesting features is that different learning models lead to sharply contrasting outcomes. Learning models that rely on best replies converge to the Cournot–Nash equilibrium, while those based on imitating the most successful player (which includes evolutionary selection) converge to the competitive (Walrasian) equilibrium (see, among others, Vega-Redondo, 1997; Vriend, 2000; Vallée and Yildizoğlu, 2009). In addition, simple and plausible trial-and-error learning converges to the collusive outcome (Huck et al., 2003, 2004a).

These theoretical predictions have led to extensive experimental analyses of these games. Huck et al. (1999, 2000), Offerman et al. (2002), and Apesteguia et al. (2010) experimentally examine "feedback effects," i.e., the impact of differences in the information provided to participants after each play of the game (feedback) on the outcomes. These studies demonstrate that outputs are higher (thus, more competitive) in settings where participants receive detailed feedback (price, aggregate quantity, quantities chosen, and profits obtained by each of the other participants in the same market) after each play of the game, compared with those settings where they receive only aggregate information as feedback (price and aggregate quantity, apart from their own choice of quantity and the resulting profits).

Note that imitating what successful others did is easier with detailed feedback than with only aggregate feedback, although, as Bigoni and Fort (2013) show, the extent to which each piece of information is used when making decisions may differ across participants. Recently, Friedman et al. (2015) showed that previously reported feedback effects are only a temporary phenomenon (observed in the first 50 periods of their 1,200-period experiment with 2 or 3 players per market), and that after a couple of hundreds of periods, even with detailed feedback, participants learn to collude. Bosch-Domènech and Vriend (2003) varied the difficulty to perform best reply (by manipulating the information required to compute payoffs in the instructions as well as the decision time) to investigate whether participants imitate successful others when undertaking the best reply is more difficult. They did not, however, observe outputs much higher in the environment more conducive to imitation.

Another well-known result from the experimental analysis of Cournot competition is "number effects," (Selten, 1973) which concerns the number of players (firms) interacting and the resulting degree of competition. In an experiment with aggregate feedback, Huck et al. (2004b) show that while some collusion can be observed when only two players interact in a market, there is near Cournot–Nash equilibrium with three players interacting, and that outcomes are never collusive and typically above the Cournot–Nash equilibrium with four or more players interacting (see Horstmann et al., 2018, for a recent survey as well as additional experimental evidence). In fact, Oechssler et al. (2016) shows that this is also the case in near-continuous 1,200-period settings similar to that of Friedman et al. (2015) and thus, unlike feedback effects, it is robust against increasing the number of periods.

In this paper, we extend the individual evolutionary learning model (first applied to Cournot games by Arifovic, 1994) by incorporating other-regarding considerations as done by Arifovic and Ledyard (2012) and apply it to Cournot games. We first fit the model to the experimental data of a repeated 3-player Cournot game (with nonlinear cost and demand functions) by Offerman et al. (2002) to fix the parameter values. We then simulate the model in n-player games (with $n \in \{2, 3, 4\}$) to make out-of-sample predictions regarding both the feedback effects and the number effects.

We then conduct our own experiment to test these out-of-sample predictions of the model. The experimental data partially confirm the prediction regarding the feedback effect; namely, the treatment with detailed feedback resulted in competitive outcomes significantly more frequently than the treatment with only the aggregate feedback for $n \in \{3, 4\}$ but not for $n = 2$. Furthermore, the prediction regarding the number effect is also partially confirmed; namely, while the model predicts the number effect to be observed only with the detailed feedback and not the aggregate feedback, it is observed with both the detailed and the aggregate feedback treatments.

The remainder of the paper is organized as follows. Section 2 presents the extended individual evolutionary learning model, along with its fit to existing experimental data and out-of-sample predictions. The experimental design to test the out-of-sample prediction of the model is presented in Section 3, followed by the results of the experiment in Section 4. Section 5 concludes.

2 Individual Evolutionary Learning Model

We extend the individual evolutionary learning (IEL) model that incorporates the other-regarding consideration proposed by Arifovic and Ledyard (2012), in which the following utility function represents each agent's preference.

$$u^i(q_i, Q_{-i}) = \pi^i(q_i, Q_{-i}) + \beta^i \bar{\pi}(Q) - \gamma^i max\{0, \bar{\pi}(q) - \pi^i(q_i, Q_{-i})\} \qquad (1)$$

where $\pi^i(q_i, Q_{-i})$ is the profit for agent i when i produces q_i and the other agents produce $Q_{-i} = \sum_{j \neq i} q_j$ in total. $\bar{\pi}(Q)$ is the average profit obtained by the agents when the total output is $Q = \sum_i q_i$. The $\beta^i \geq 0$ captures the degree to which i cares about the average profit. Because $\beta^i \geq 0$, the higher the average profit is, the higher the i's utility is. The term $max\{0, \bar{\pi}(q) - \pi^i(q_i, Q_{-i})\}$ captures how much less profit i makes compared with the average profit, and $\gamma^i \geq 0$ the degree to which i cares about making less profit than the average. As we can see from the negative sign in front of the third term, the less profit i makes compared with the average, the lower i's utility is.

The profit for firm i is defined as follows.

$$\pi^i(q_i, Q_{-i}) = (a - b\sqrt{Q})q_i - cq_i^d \quad (2)$$

In this paper, we assume $\{a, b, c, d\} = \{45, \sqrt{3}, 1, 3/2\}$ because these are the values used in the experiment of Offerman et al. (2002).

The existence of other-regarding preferences is often posited as an explanation for seemingly irrational behavior in experimental settings. For example, in the context of Cournot games, Iriş and Santos-Pinto (2014) have used other-regarding preferences to explain several regular deviations from theoretical predictions.

As the superscript i on β^i and γ^i indicates, we assume agents are heterogeneous. Andreoni and Miller (2002) report that the vast majority of behavior in a standard dictator game can be explained by the presence of a range of heterogeneous other-regarding preferences. Blanco et al. (2011) find further evidence for a range of heterogeneous other-regarding preferences among the experimental population. However, they also find that these preferences are unstable or manifest differently in different experimental settings.

Furthermore, as Bigoni and Fort (2013) show in their experiment where participants can choose which information to consider, not all the participants utilize the same feedback information, which may be a result of participants belonging to different types, as Rassenti et al. (2000) suggest. We, therefore, assume that there are the following five types of agents in the population:

- Individualistic: $\beta^i = 0$, $\gamma^i = 0$

- Cooperative: $\beta^i \sim U(0, B)$, $\gamma^i = 0$ where B is the upper bound

- Retaliatory: $\beta^i = 0$, $\gamma^i \sim U(0, G)$ where G is the upper bound

- Mixed: $\beta^i \sim U(0, B)$,$\gamma^i \sim U(0, G)$

- Flexible: A flexible agent starts as an Individualistic agent. At the start of each turn, the agent updates their γ^i_t (β^i_t), entering the utility function in equation (a), with a fixed probability ρ as follows:

$$\gamma^i_t = \frac{\bar{\pi}^{-i}_{t-1} - \pi^i_{t-1}}{\bar{\pi}^{-i}_{t-1}}(G - \gamma^i_{t-1}) \times \omega_f + \gamma^i_{t-1} \text{ if } \bar{\pi}^{-i}_{t-1} \geq \pi^i_{t-1}$$

$$\beta^i_t = \frac{\pi^i_{t-1} - \bar{\pi}^{-i}_{t-1}}{\pi^i_{t-1}}(B - \beta^i_{t-1}) \times \omega_f + \beta^i_{t-1} \text{ if } \bar{\pi}^{-i}_{t-1} < \pi^i_{t-1}$$

  where $\bar{\pi}^{-i}_{t-1}$ is the average profit of other agents in the previous turn, and ω_f measures how sensitive the Flexible type agent is to others' profits.

2.1 Dynamics of IEL

In IEL, each agent i, regardless of type, has a set of J actions (quantities), S^i_t with $|S^i_t| = J$, to choose from in each period t. The initial set of actions, S^i_0, is randomly generated from the set of possible actions for player i, $\mathcal{S}^i$, and the initial action, q^i_0, is randomly selected from S^i_0. Learning in IEL takes the form of updating S^i_t at the end of each period by the following two steps: Experimentation and Replication.

Experimentation introduces new actions into S^i_t that otherwise might never have a chance to be tried. This ensures that a certain amount of diversity is maintained in S^i_t. For each $q^i_{j,t} \in S^i_t$ ($j = 1, ..., J$), with probability μ, a new action, q, is selected randomly according to $N(q^i_{j,t}, \sigma|\mathcal{S}^i)$ and replaces it. This means that the 'new' action that is replacing $q^i_{j,t}$, on average, has the same value. μ and σ are parameters of the model. Let $\widetilde{S}^i_t$ be the set of actions after experimentation.

Replication reinforces, after experimentation, actions that would have been good choices in period t. This allows potentially better-performing actions to replace worse-performing actions. We assume that the performance of $q^i_{j,t} \in \widetilde{S}^i_t$ is evaluated according to the type-dependent utility function stated above, conditional on the information i received regarding the outcome in period t, $I^i(x_t)$. Let $v^i(q^i_{j,t}|I^i(x_t)) = u^i(q^i_{j,t}, Q_{-i}|I^i(x_t))$ be the computed performance level of $q^i_{j,t}$ holding constant the output levels of other firms. This measures the utility i that would have been obtained had i played $q^i_{j,t}$ in period t holding everything else constant. The element of the new action set S^i_{t+1}, $q^i_{j,t+1}$, is chosen as follows. For each $j = 1, \cdots, J$, select two members of $\widetilde{S}^i_t$, $q^i_{m,t}$and $q^i_{l,t}$, uniformly randomly with replacement. Then,

$$q^i_{j,t+1} = \begin{cases} q^i_{m,t} \text{ with probability } \frac{e^{v^i(q^i_{m,t}|I^i(x_t))}}{e^{v^i(q^i_{m,t}|I^i(x_t))} + e^{v^i(q^i_{l,t}|I^i(x_t))}} \\ q^i_{l,t} \text{ otherwise.} \end{cases} \tag{3}$$

Once S^i_{t+1} is generated, in period $t+1$, the action of agent i is chosen probabilistically according to

$$\psi^i_{k,t+1} = \frac{e^{\lambda v^i(q^i_{k,t+1}|I^i(x_t))}}{\sum_{j=1}^J e^{\lambda v^i(q^i_{j,t+1}|I^i(x_t))}} \tag{4}$$

where $\lambda \geq 0$ is a parameter of the model that governs the sensitivity of action choice to their performance level.

2.2 Model fitting

To fix values of the parameters of the IEL model, we fit the model to the data reported in Offerman et al. (2002). In Offerman et al. (2002), participants play a 3-player repeated Cournot game as specified above for 100 periods. Three treatments, Q, Qq, and $Qq\pi$, are considered by varying the information participants receive after each play of the game (feedback after each period) across treatments. In Treatment Q, participants are informed of the realized price and aggregate quantity (in addition to their own quantity and profit). In Treatment Qq, in addition to the information provided in treatment Q, participants are informed of the quantity chosen by each of the other participants in the group. In Treatment $Qq\pi$, participants are informed, in addition to those provided in Qq, profits obtained by each of the other participants in the group.

These variations in feedback after each period across the three treatments influence the simulations of our model as follows.

In Treatment Q , agents only know the aggregate output. They cannot calculate the average profit without information on the individual output level because the cost function is nonlinear. Therefore, we assume that all agents are Individualistic agents in our simulation. This restriction means that we have three parameters—μ, σ, and λ—in our model. We have conducted a grid search in the space of $\mu \in \{0.03, 0.07, 0.1, 0.3, 0.6\}$, $\sigma \in \{0.3, 0.5, 0.7, 2, 5, 10, 20\}$, and $\lambda \in \{6 \times 10^{-4}, 8 \times 10^{-4}, \cdots, 24 \times 10^{-4}\}$ (in step size of 2×10^{-4}). For each combination of parameter values, we run 30 simulations. Based on the measure of fit described in Appendix A, we select $\mu = 0.03$, $\sigma = 2$, and $\lambda = 0.002$ as the best parameter values.

In Treatment Qq **and** $Q q\pi$, in addition to the aggregate output, agents also know the individual output of others in the group in Qq. In $Qq\pi$, they are further informed of the profits of others in the group. Although it is hard to say how well agents compute others' profits from their outputs given the nonlinearity of the cost function, it is possible. Therefore, we simulate the model with the full set of parameters. We vary the proportions of the five types of agents in the population (f_I, f_C, f_R, f_F, and f_M for Individualistic, Cooperative, Retaliatory, Flexible, and Mixed, respectively, while maintaining $\sum_x f_x = 1$ where $x \in I, C, R, F, M$), B (the upper bound of β^i), G (the upper bound of γ^i), ρ and ω^f (probability of updating and sensibility to others' profit for Flexible type), and fix the values of μ, σ, and λ to those we selected for Treatment Q.

Because of the large parameter space that needs to be investigated, we conduct a grid search in two steps, as detailed in Appendix B: First, a coarse grid search, and second, a finer grid search for the subregion of the parameter space that resulted in the highest fit under the coarse grid search. For each combination of parameter values considered, 50 simulations are run. We find, in addition to $\mu = 0.03$, $\sigma = 2$, and $\lambda = 0.002$ that were fixed, the sets of parameter values reported in Table 1 to be the best ones in fitting the model to the experimental data according to the criterion reported in Appendix A.[1]

Comparing the best set of parameter values for treatment Qq and $Qq\pi$ reported in Table 1 reveals an interesting difference. Our simulations show that Qq can be best described by the interactions between Individualistic and Cooperative types, although more than two-thirds of the agents are the former. In contrast, for $Qq\pi$, there is no Individualistic type. Instead, it is best described by the interactions between

[1]There was another set of parameter values for Qq that was equally good, but we have used the one reported in Table 1 for our out-of-sample simulation to be reported in the next subsection. See Appendix B.

Table 1: The best set of parameter values for the three treatments.

	μ	σ	λ	f_I	f_C	f_R	f_F	f_M	B	G	ρ	ω^f
Q	0.03	2.0	0.002	1.0	-	-	-	-	-	-	-	-
Qq	0.03	2.0	0.002	0.8	0.2	0.0	0.0	0.0	80	0	0	0
$Qq\pi$	0.03	2.0	0.002	0.0	0.2	0.4	0.4	0.0	60	50	0.1	1.0

Cooperative, Retaliatory, and Flexible, with the Cooperative share being the lowest. This difference in composition, particularly the prevalence of the Individualistic type in Treatment Qq, may reflect the difficulties participants (in the experiment) have in computing others' profits from the quantity information.

Figure 1 depicts the distributions of group outputs in three treatments from the experiment of Offerman et al. (2002) (Top) and the IEL model (Bottom). The model does a good job of capturing the distributions from the experiment in Treatment Q that is centered around the Cournot–Nash equilibrium. While the model does not exactly replicate the distributions from the experiment in Treatments Qq and $Qq\pi$, it does capture the qualitative differences across treatments; namely, in Treatment $Qq\pi$, collusive and competitive outcomes are much more frequently observed compared with treatments Qq and Q. Furthermore, in Treatment Qq, the output distribution is skewed toward collusive ones compared with Treatment Q.

Thus, the simulation of the fitted model demonstrates the feedback effect. However, this is because the model is fitted to the data that demonstrate the effect. We conduct out-of-sample simulations by varying the number of players in the next subsection to generate a set of hypotheses to be tested using our experiments.

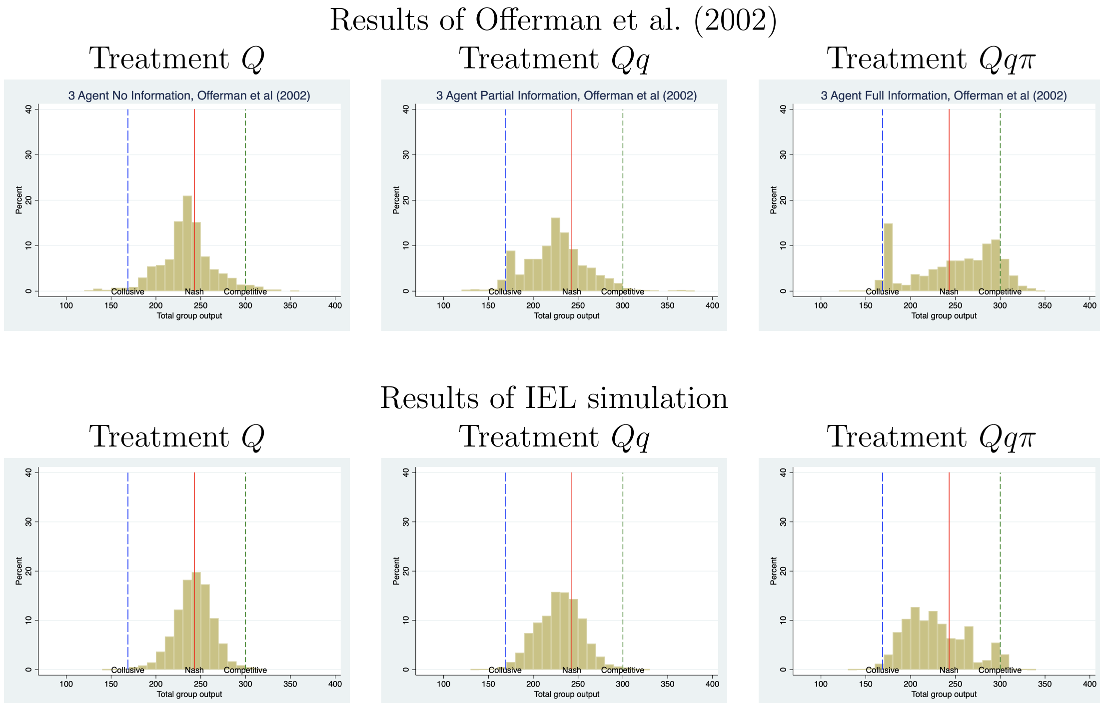

Figure 1: TOP: Distributions of group outputs in the three treatments of Offerman et al. (2002). Regenerated from the original experimental data. BOTTOM: Distributions of 15,000 simulated group outputs (50 simulations by 100 periods in 3 treatments) from the best set of parameter values.

2.3 Out-of-sample simulations: Varying number of players

Figure 2 presents the results of simulations in which we vary the number of players, $n \in \{2, 3, 4\}$, in each of the three information treatments, Q, Qq, and $Qq\pi$.[2] The $n = 3$ case is identical to that shown in Figure 1 but replicated to facilitate the comparison.

Figure 2 shows that in Treatment Q, the outcomes are distributed around Cournot–Nash regardless of the number of players. Thus, the model predicts no number effect

[2]Participants in our experiment (Section 3) could choose from integer values between 40 and 170 in 2-player treatments, 40 and 113 in 3-player treatments, and 40 and 85 in 4-player treatments, inclusive, as a quantity in each period. In simulating the model, we introduced the same constraints in quantities.

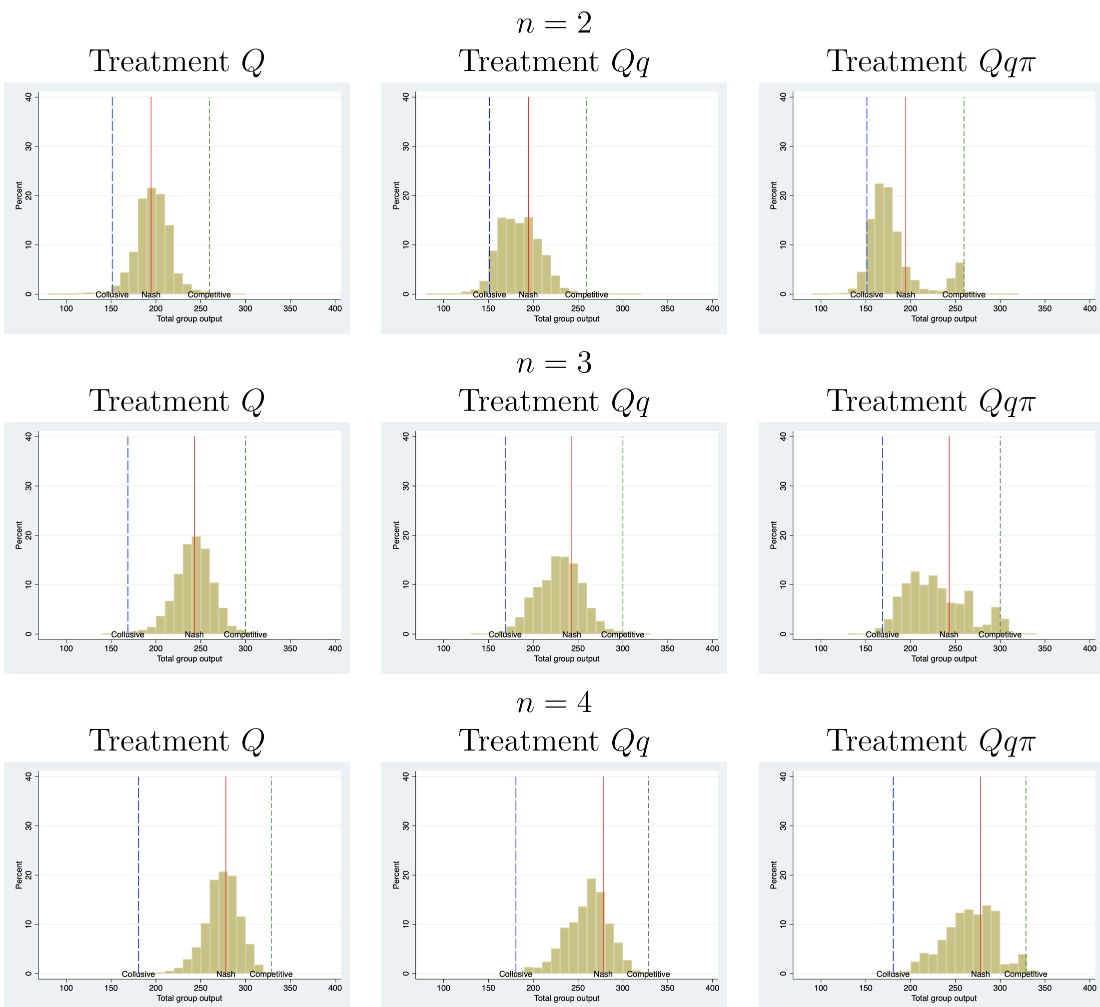

Figure 2: Distribution of the 45,000 simulated group outputs (50 simulations by 100 periods in 9 treatments) from the best set of parameter values.

under Treatment Q.[3] In Treatment $Qq\pi$, on the other hand, when $n = 2$, collusive and competitive outcomes are more likely to be observed than Cournot–Nash outcomes. For $n = 3$, the collusive outcome becomes less likely (compared with $n = 2$), while the Cournot–Nash and collusive outcomes instead become more prevalent. The output distribution is centered around the Cournot–Nash for $n = 4$, thereby demonstrating the number effect.

These differences in the simulated outcomes between treatment Q and $Qq\pi$ constitute our main hypotheses to be tested in the experiment.

Hypothesis 1 *The feedback effects (the more detailed the feedback is, the more competitive outcomes become) exist. Furthermore, when the number of players is 2 or 3, it is not just that there are more competitive outcomes, but also that we observe more collusive outcomes, with more detailed than aggregate feedback.*

Hypothesis 2 *The number effect (the larger the number of players, the less collusive the outcome) arises with detailed feedback (Treatment $Qq\pi$) but not with aggregate feedback (Treatment Q).*

Note that while Horstmann et al. (2018) reports the number effect with detailed feedback, which is consistent with our Hypothesis 2, Huck et al. (2004b) details the number effect with aggregate feedback, which is inconsistent with our second hypothesis. However, there are several differences in setting that make comparison of our simulation results and these existing experimental results difficult. First, and most importantly, these two experiments are based on linear demand and cost functions (in the case of Horstmann et al. (2018) with zero cost), while our simulation is based on nonlinear cost and demand functions. Second, there are 25 and 60 periods in

[3] However, the relative locations of the collusive, Cournot–Nash, and competitive outcomes change depending on the number of players. Namely, Cournot–Nash is closer to collusive than to competitive in $n = 2$, while the opposite is the case for $n = 4$.

the experiments in Huck et al. (2004b) and Horstmann et al. (2018), respectively, while our simulation is based on 100 periods. The results in Huck et al. (2004b), which are based on 25 periods, therefore can arise because of insufficient learning by participants. Because we are unaware of experimental analysis concerning the number effect in the Cournot game using the nonlinear demand and cost functions considered by Offerman et al. (2002) under varying feedback conditions, we have conducted a new set of experiments of our own to test these hypotheses.

3 Experiment Design

Table 2 presents the treatments in our two-by-three design: we vary feedback information (Partial vs. Full) and the number of players in a market ($n \in \{2, 3, 4\}$). In Partial-Info treatments, the players received feedback information on the total quantity, the price, and their own individual output and profit at the end of each period (thus, identical to treatment Q above). In Full-Info treatments, the players received additional information on the individual outputs and profits of others in the same group (thus, identical to treatment $Qq\pi$). When they made decisions, they could always review the full history of the individual outputs in a graph and a table (see Appendix D.3). Both information treatments were run in 2-, 3-, and 4-player Cournot games.

We repeated the game for 100 periods. In each period, firms must decide simultaneously how much to produce. They can only choose integer values between 40 and 170 in 2-player treatments, 40 and 113 in 3-player treatments, and 40 and 85 in 4-player treatments, inclusive. Note all firms face the same cost and demand functions. All firms in the industry receive the same price for each commodity produced. Both the relationship between own production and costs and aggregate production and

Table 2: Summary of treatments.

	Feedback information	
Number of players	Partial-Info (π_i,Q,P)	Full-Info(π_i,Q,P,q_j,π_j)
2 players	Partial2	Full2
3 players	Partial3	Full3
4 players	Partial4	Full4

Notes: All the feedback information concerns the previous period. Subscript i refers to the firm itself; subscript j concern other firms in the industry; π denotes profit; Q denotes aggregate output; and P denotes price.

price were illustrated in interactive graphs. Players needed to correctly answer questions about the cost and demand functions before proceeding with the experiment. The introduction, along with the interactive graphs, were made available throughout the experiment (see Appendix D.2 for an example of the on-screen instruction.)

Unlike Offerman et al. (2002) but as in Bigoni and Fort (2013), we provided an on-screen profit calculator to help players figure out their best responses. They could enter their own hypothetical output and the aggregate output of others in the market to get their own hypothetical profit. Players used the calculator on average three times per period, with 72 percent of players using it more than once per period.

During the experiment, players earned experimental points according to the profits they obtained. Participants were rewarded based on the sum of the points earned during 100 periods. At the end of the experiment, the experimental points were exchanged for Canadian dollars at an exchange rate of 5,000 experimental points = 1 Canadian dollar.

The experiment is programmed in oTree (Chen et al., 2016). The Cournot game was preceded by a part of the advanced version of Raven's Progressive Matrix test (Raven, 1998). For the Raven test, we selected 16 questions from the original test

Table 3: Summary of the number of groups and average earnings (in Canadian dollars).

Treatment	Num. of Groups	Cournot game	Raven test	Total earning
Partial2	6	$20.7	$2.7	$30.4
Partial3	6	$14.3	$2.3	$23.6
Partial4	6	$11.0	$2.5	$20.5
Full2	9	$19.4	$2.6	$29.0
Full3	7	$12.7	$3.1	$22.8
Full4	6	$9.8	$2.5	$19.3

Notes: Total earnings equal to the sum of the earnings in Cournot game, Raven test, and $7 participation fee.

to be finished within 10 minutes.[4] In each question, players analyzed a geometric pattern and identified a missing part to complete the series. For each question answered correctly, players were paid 0.25 Canadian dollars (see Appendix D.1 for the Raven test instructions). The Raven test measures "fluid intelligence"; that is, "the capacity to think logically, analyze and solve novel problems, independent of background knowledge" (Mullainathan and Shafir, 2013, p.48). Recently, Proto et al. (2019) reported that participants with higher scores on the Raven test are more likely to achieve and sustain cooperation in an infinitely repeated prisoners' dilemma game than those with lower scores. Because our experiment on Cournot games involves similar trade-offs between short-run temptation and long-run gain, we decided to gather this information to control its effect, if any, when analyzing the data.

Laboratory experiments were conducted between September 2019 and January 2020. Participants were recruited at Simon Fraser University. A total of 117 subjects participated. An experiment lasted between 1.5 and 2 hours. Average earnings were 23.5 Canadian dollars (see Table 3 for the number of groups and earnings under different treatments). The number of participants differs across treatments because

[4]The original test consists of 48 questions to be answered in 30 to 40 minutes. We selected 1 in every 3 questions while retaining the order of questions.

of variation in the rate of show up or appearance.

4 Results of the Experiment

Figure 3 shows the distribution of group outputs observed for each of our six treatments. In the three Partial-Info treatments, total group outputs are distributed around the Cournot–Nash equilibrium, just as we observed in the simulations of the extended IEL shown in Figure 2. In contrast, in the three Full-Info treatments, more competitive group outputs are observed compared with the Partial-Info treatment for each group size. While the frequency of collusive outcomes is low for the Full2 treatment, it is much higher in Full3 than in the Full4 treatment.

Let us define “collusive outcomes” to be outcomes such that $Q_g \leq \frac{1}{2}\left(Q_{Nash} + Q_{Coll}\right)$ and “competitive outcomes” to be outcomes such that $Q_g \geq \frac{1}{2}\left(Q_{Nash} + Q_{Comp}\right)$ where Q_g is the total output of firms, and Q_{Nash}, Q_{Comp}, and Q_{Coll} are the corresponding Nash, competitive, and collusive levels of total outputs, respectively. We also define remaining outcomes that are around Q_{Nash} to be “Nash outcomes.”

Figure 4 shows the distribution of the relative frequencies of “Nash”, “Collusive”, and “Competitive” outcomes for each group across the various treatments. In each plot, filled marks represent groups in the Partial-Info treatment while unfilled marks represent the Full-Info treatment. The point at the very top of the triangle corresponds to $(``Nash", ``Collusive", ``Competitive") = (100, 0, 0)$; that is, the group's output was Nash outcomes 100% of the time. The left and the right bottom apexes of the triangle corresponds to $(``Nash", ``Collusive", ``Competitive") = (0, 100, 0)$ and $(0, 0, 100)$, respectively. That is, the group's output was collusive or competitive outcomes 100% of the time. The p values beneath the plot are for the Mann–Whitney test (two-tailed) comparing the frequencies of collusive and competitive outcomes

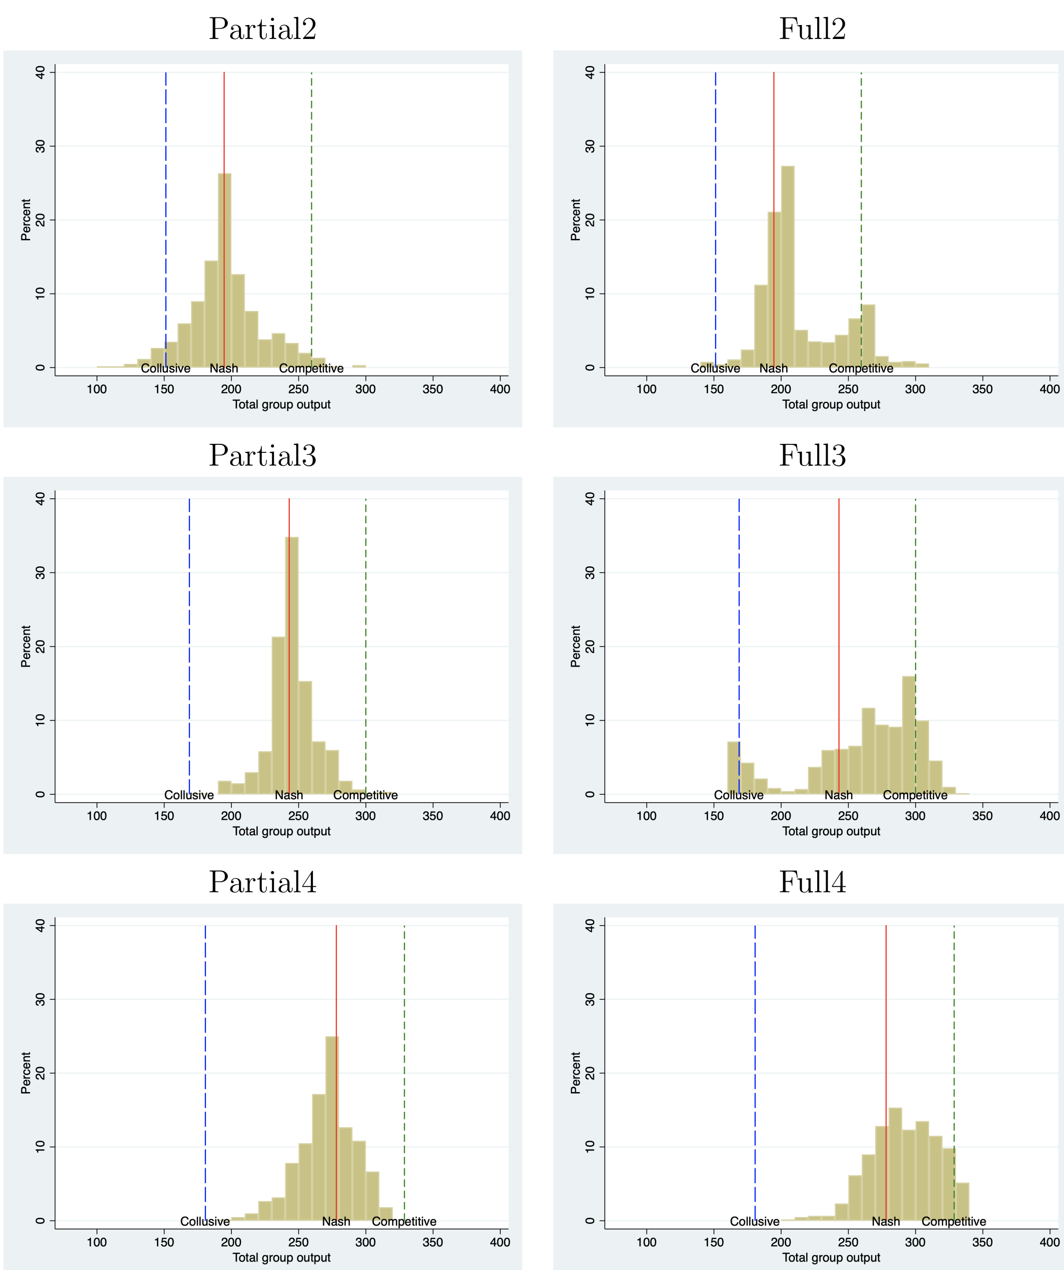

Figure 3: Distribution of group outputs across the six treatments.

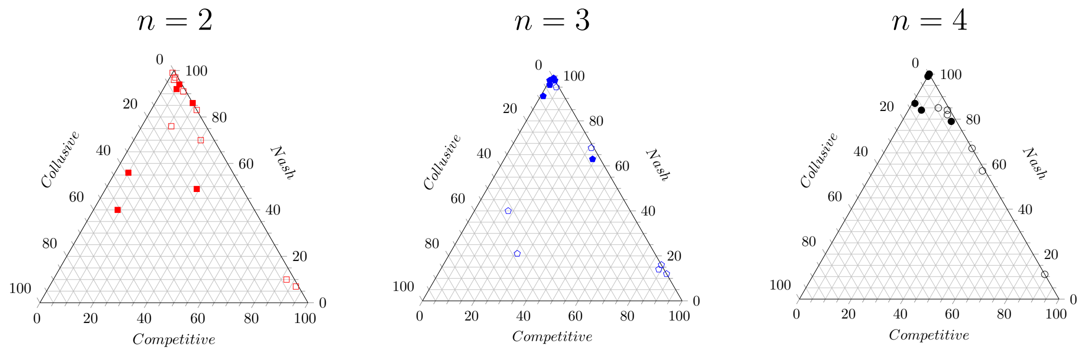

p=0.23 (Collusive) p=0.81 (Competitive) | p=0.94 (Collusive) p=0.01 (Competitive) | p=0.56 (Collusive) p=0.02 (Competitive)

Figure 4: Distribution of the frequencies of "Nash", "Collusive", and "Competitive" outcomes. Filled (unfilled) marks represent partial (full) info treatment. The p values beneath the plot are from a Mann–Whitney test (two-tailed) comparing the frequencies of collusive and competitive outcomes between the Partial and Full-Info treatments.

between the Partial-Info and Full-Info treatments.

For $n = 3$ (middle panel), we can see that while the filled dots are mostly distributed around the top apex (which means the outcomes are mostly "Nash" in most of the groups), the unfilled dots are distributed either toward the bottom left or right apexes of the triangle (which means that "Collusive" or "Competitive" outcomes are more frequently observed than "Nash"). In particular, we observe that the frequencies of competitive outcomes are significantly higher under the Full-Info than the Partial-Info treatment for $n \in \{3, 4\}$. We therefore make the following observation.

Observation 1 *The feedback effect exists in the experiment. Full-Info treatments result in competitive outcomes significantly more frequently than the Partial-Info treatments for $n \in \{3, 4\}$.*

Observation 1 is consistent with the first part of Hypothesis 1. However, the second part of Hypothesis 1 (about collusive outcomes) is not confirmed in our data.

Table 4: Tacit collusion.

Dependent variable: Tacit Collusion	
	(1)
	none
FID	-1.231***
	(0.013)
Agent Number	-0.666***
	(0.002)
Agent Number *times* FID	0.313***
	(0.004)
Constant	5.352***
	(0.005)
Group FE	Yes
Observations	4000

standard deviations in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: FID is a dummy variable: 1 if under Full-Info treatments, 0 otherwise.

We now investigate our Hypothesis 2 about the number effect. We do not, however, observe any significant differences across the three group sizes in terms of the frequency of collusive outcomes, neither in the Partial-Info (p=0.40, Kruskal–Wallis test) nor Full-Info treatment (p=0.71, Kruskal–Wallis test).

To check the number effect more closely, following Horstmann et al. (2018), Engel (2007), and Suetens and Potters (2007), we measure the degree of tacit collusion by $\phi_t^E \equiv \frac{\overline{p_t} - p^{Nash}}{p^{JPM} - p^{Nash}}$ where $\overline{p_t}$ is the market price.[5] We then run a group fixed effects regression to test the impact of the number of agents on this measure.[6] Table 4 shows the result.

Note that the degree of tacit collusion is lower in the Full-Info treatment as demonstrated by the negative and significant coefficient for the Full-Info treatment dummy (FID). This is consistent with our analyses above on the feedback effect summarized

[5] Replacing p^{Nash} by p^{Walras} yields similar results.

[6] We use group fixed effects regression because there is substantial heterogeneity in terms of the distribution of outcomes.

in Observation 1. The negative and significant coefficient of *Agent Number* in Table 4 indicates a lower degree of collusion for larger markets, demonstrating the existence of the number effect in the Partial-Info treatment which is contrary to our Hypothesis 2. In fact, the positive and significant coefficient of *Agent Number* $\times$ *FID* shows that the number effect is weaker under the Full-Info treatment compared with the Partial-Info treatment. We, therefore, make the following observation.

Observation 2 *The number effect exists in both the Full and Partial-Info treatments. The number effect is weaker, however, under the Full-Info treatment, contrary to Hypothesis 2.*

4.1 Group behavior

Why are collusive outcomes rare in the Full2 treatment compared to its $n = 3$ counterpart, and unlike what the simulated IEL predicts? One possible reason is the differences in average cognitive ability among participants across treatments. While we have randomized participants across treatments, it is possible that the distribution of participants in terms of their cognitive ability was not the same across the treatments. Because it has been shown that participants with higher cognitive ability tend to achieve a higher degree of cooperation (Proto et al., 2019), variation in cognitive ability across sessions (or more importantly, across treatments) can result in participants achieving more cooperative (collusive) outcomes in some sessions (treatments) than in others.

As a first check, we run the following simple linear regression:

$$Average\ Output_{ij} = \beta Average\ Raven\ score_{ij} + Treatment_j + \epsilon_{ij} \tag{5}$$

where $Average\ Output_{ij}$ and $Average\ Raven\ score_{ij}$ are the average individual out-

Table 5: Group output and the average score of the Raven test.

Dependent variable: Average Individual Output	(1) Full Sample	(2) Partial-Info	(3) Full-Info
Average Raven Score	-1.635*	-0.461	-2.996*
	(0.873)	(0.711)	(1.540)
Treatment FE	Yes	Yes	Yes
Observations	40	18	22

standard deviations in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Column 1 shows the result from the full sample. Column 2 uses the subsample of three Partial-Info treatments. Column 3 uses the subsample of three Full-Info treatments.

puts and the average score of the Raven test for participants in group i with treatment j, respectively, and $Treatment_j$ is the treatment fixed effect.

Table 5 details the results. The negative coefficient of *Average Raven Score* in Column 1 means that groups with a higher average score on the Raven test tend to produce less and are more likely to reach the collusive outcome. On average, a one-point increase in the average Raven score leads to a 1.63 decrease in average individual output. This finding is consistent with Proto et al. (2019). In our experiment, this significant correlation between the Raven test score and the average output is driven by the groups in the Full-Info treatments, as shown in Column 3. This correlation is not significant in groups under the Partial-Info treatments.

Figure 5 compares the average group Raven scores from all the treatments. The average Raven score for the Full3 treatment is much higher than in the other treatments. This might indeed be the reason we observe more collusive outcomes in the Full3 treatment than in the Full2 treatment, contrary to what our simulation predicted.

To better understand the mechanism of the observed feedback effect in our exper-

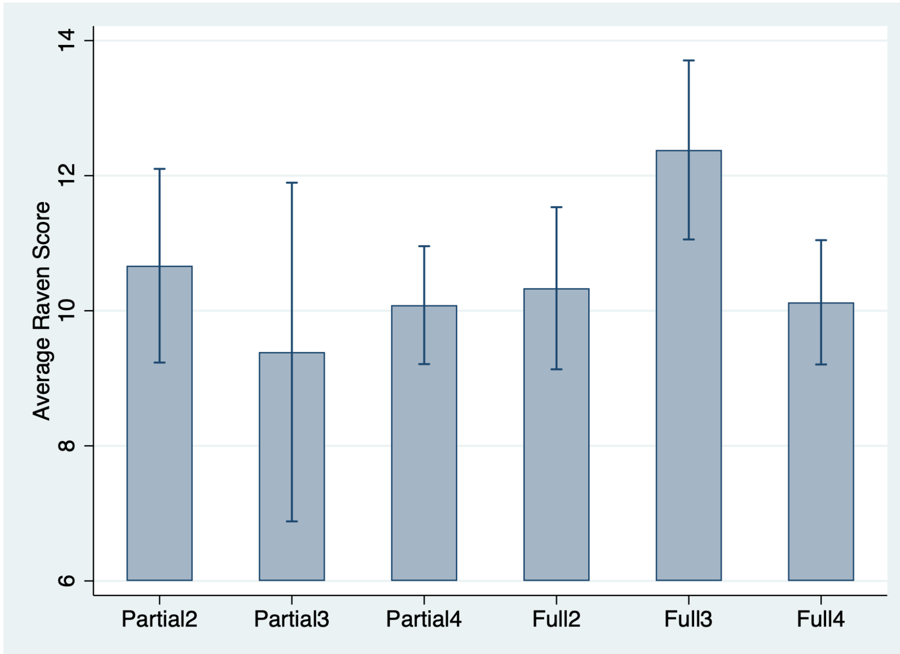

Figure 5: The average group Raven score across treatments.

iment, we analyze the data at the individual level in the next subsection.

4.2 Individual behavior

At the individual level, participants exhibit a large degree of heterogeneity. To quantify their behavior, we divide the output levels into three categories: Individualistic, Cooperative, and Retaliatory. We define an individual output decision to be Cooperative if the output is 10 units lower than the best response output level based on the total output of others in the previous round; Retaliatory if the output is 10 units higher than the best response output level based on the total output of others in the previous round; and Individualistic otherwise. In addition to the above three types, we also define an action to be Flexible if the agent increases (or decreases) their output by at least 10 units, immediately following an increase (or decrease) in the aggregate output of the others in the group in the previous round.

We use these definitions because the total output of others is available in all treatments, and participants can easily look for their profit-maximizing output level, conditional on others' total outputs, by using the calculators provided in the experiment.[7] In fact, an average agent uses the profit calculator 313 times throughout the experiment.[8]

Figure 6 illustrates the composition of the various decision types for the two information treatments. The participants under Full-Info treatments make more Retaliatory decisions and fewer Individualistic decisions than those under Partial-Info treatments. Given aggregate output levels are provided to the participants in all treatments, the significantly higher percentage of Retaliatory decisions under the Full-Info treatments must be driven by the extra information (namely, individual quantities and profits) provided.

To understand the effect of various action types on average output, the following regression is conducted:

$$Average\ Output_{ij} = \beta_1 Cooperative_{ij} + \beta_2 Retaliatory_{ij} + \beta_3 Flexible_{ij} + Group_j + \epsilon_{ij} \tag{6}$$

where $Average\ Output_{ij}$ is the average output per round for individual i in the group j; $Cooperative_{ij}$, $Retaliatory_{ij}$, and $Flexible_{ij}$ are the number of Cooperative, Retaliatory, and Flexible decisions (across 100 periods), respectively, for individual i in group j; $Group_j$ is the group fixed effect; and ϵ_{ij} is the error term.[9]

Figure 7 provides the estimated coefficients of the various decision types. Not surprisingly, an increase in the number of Cooperative decisions significantly decreases

[7] Although the information available is the total group output, participants can easily deduct their own output from the total group output to compute the total output of other group members.

[8] 81 percent of participants use it at least 50 times throughout the experiment.

[9] Individualistic decisions are dropped given collinearity.

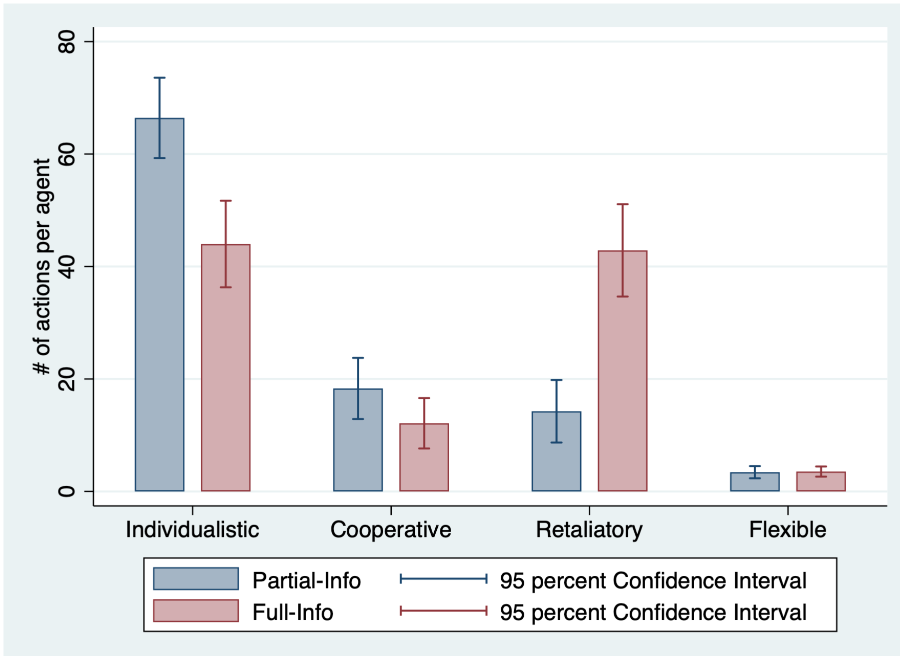

Figure 6: Number of various types of individual action by information type.

average individual output, while an increase in the number of Retaliatory decisions significantly increases average individual output. Flexible decisions have no significant effect on average individual output. Importantly, these effects are similar across the different information treatments and not statistically significantly different at the 5% level (see, Table 6, in Appendix C), suggesting that the different distributions of outputs across treatments we saw in Figure 3 are not driven by the relative strengths of the various types of actions.

5 Summary and Discussion

In this paper, we extend the individual evolutionary learning model by incorporating other-regarding considerations and apply the model to some Cournot games. Based on the model fitted to the experimental data of a repeated 3-player Cournot game (with nonlinear cost and demand functions), we make out-of-sample predictions regarding

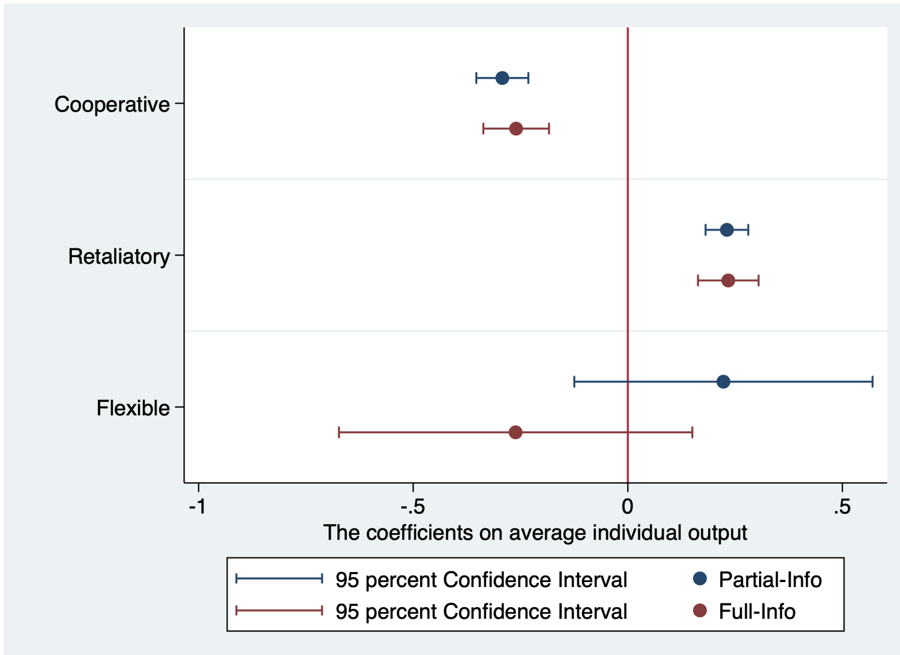

Figure 7: The estimated coefficients of various decision types for the two treatments.

the feedback and number effects and test them using the data gathered via newly conducted experiments. The prediction regarding the feedback effect is partially confirmed. Namely, we observe it in the 3- and 4-player games but not in the 2-player game. The prediction regarding the number effect is also partially confirmed in that while the model predicts the number effect to be observed with detailed feedback, and not under aggregate feedback, the effect is observed with both types of feedback.

We investigate the role of cognitive ability and individual behavior to better understand the experimental outcomes. The data suggest that the differences in the average cognitive ability of participants across treatments may account for a higher frequency of collusive outcomes observed in the 3-player game with detailed feedback than the 2-player game with detailed feedback, although the model predicts the opposite.

The individual level analyses reveal that participants are making Retaliatory decisions (that is to produce much more than the myopic best response level) much more frequently with detailed feedback than with aggregate feedback, which results in more competitive outcomes in the former than the latter.

We believe the approach taken in the paper—namely, to calibrate the parameter values of a boundedly rational behavioral model by fitting the model to an existing experimental data set and then testing the out-of-sample predictions of the calibrated model with newly corrected data—can be a powerful tool for advancing research in behavioral and experimental economics. A well-known problem in working with boundedly rational behavioral models is the so-called "wilderness" of bounded rationality (Sims, 1980; Hommes, 2013); that is, because of the large degree of freedom modelers have in constructing these models, systematic investigation becomes extremely difficult. By restricting our attention to those models with a solid empirical (experimental) foundation, and further testing their out-of-sample prediction with newly gathered data, we can narrow down the set of models to advance one's investigation.

References

Andreoni, J. and J. Miller (2002): "Giving according to GARP: An experimental test of the consistency of preferences for altruism," *Econometrica*, 70, 737–753.

Apesteguia, J., S. Huck, J. Oechssler, and S. Weidenholzer (2010): "Imitation and the evolution of Walrasian behavior: Theoretically fragile but behaviorally robust," *Journal of Economic Theory*, 145, 1603–1617.

Arifovic, J. (1994): "Genetic algorithm learning and the cobweb model," *Journal of Economic Dynamics and Control*, 18, 3–28.

Arifovic, J. and J. Ledyard (2012): "Individual evolutionary learning, other-regarding preferences, and the voluntary contributions mechanism," *Journal of Public Economics*, 96, 808–823.

Bigoni, M. and M. Fort (2013): "Information and learning in oligopoly: An experiment," *Games and Economic Behavior*, 81, 192–214.

Blanco, M., D. Engelmann, and H. T. Normann (2011): "A within-subject analysis of other-regarding preferences," *Games and Economic Behavior*, 72, 321–328.

Bosch-Domènech, A. and N. J. Vriend (2003): "Imitation of Successful Behaviour in Cournot Market," *Economic Journal*, 113, 495–524.

Chen, D. L., M. Schonger, and C. Wickens (2016): "oTree – An open-source platform for laboratory, online, and field experiments," *Journal of Behavioral and Experimental Finance*, 9, 88–97.

Engel, C. (2007): "How much collusion? A meta-analysis of oligopoly experiments," *Journal of Competition Law and Economics*, 4, 491–549.

Friedman, D., S. Huck, R. Oprea, and S. Weidenholzer (2015): "From imitation to collusion: Long-run learning in a low-information environment," *Journal of Economic Theory*, 155, 185–205.

Hommes, C. (2013): *Behavioral Rationality and Heterogeneous Expectations in Complex Economic Systems*, Cambridge, UK: Cambridge University Press.

Horstmann, N., J. Krämer, and D. Schnurr (2018): "Number effects and tacit collusion in experimental oligopolies," *Journal of Industrial Economics*, LXVI, 650–700.

Huck, S., H.-T. Normann, and J. Oechssler (1999): "Learning in Cournot Oligopoly - An Experiment," *Economic Journal*, 109, C80–C95.

——— (2000): "Does information about competitors' actions increase or decrease competition in experimental oligopoly markets?" *International Journal of Industrial Organization*, 18, 39–57.

——— (2003): "Zero-knowledge cooperation in dilemma games," *Journal of Theoretical Biology*, 220, 47–54.

——— (2004a): "Through trial and error to collusion," *International Economic Review*, 45, 205–224.

——— (2004b): "Two are few and four are many: number effects in experimental oligopolies," *Journal of Economic Behavior and Organization*, 53, 435–446.

Iriş, D. and L. Santos-Pinto (2014): "Experimental Cournot oligopoly and inequity aversion," *Theory and Decision*, 76, 31–45.

Mullainathan, S. and E. Shafir (2013): *Scarcity: Why Having Too Little Means So Much*, New York, NY: Times Books, Henry Holt and Company, LLC.

Oechssler, J., A. Roomets, and S. Roth (2016): "From imitation to collusion: a replication," *Journal of Economic Science Association*, 2, 13–21.

Offerman, T., J. Potters, and J. Sonnemans (2002): "Imitation and belief learning in an oligopoly experiment," *Review of Economic Studies*, 69, 973–997.

Proto, E., A. Rustichini, and A. Sofianos (2019): "Intelligence, Personality and Gains from Cooperation in Repeated Interactions," *Journal of Political Economy*, 127, 1351–1390.

Rassenti, S., S. S. Reynolds, V. L. Smith, and F. Szidarovsky (2000): “Adaptation and convergence of behavior in repeated experimental Counot games,” *Journal of Economic Behavior and Organization*, 41, 117–146.

Raven, J. C. (1998): *Raven's Advanced Progressive Matrices (APM)*, San Antonio, TX: Pearson, 2003 ed.

Selten, R. (1973): “A simple model of imperfect competition, where 4 are few and 6 are many,” *International Journal of Game Theory*, 2, 141–201.

Sims, C. A. (1980): “Macroeconomics and Reality,” *Econometrica*, 48, 1–48.

Suetens, S. and J. Potters (2007): “Bertrand colludes more than Cournot,” *Experimental Economics*, 10, 71–77.

Vallée, T. and M. Yildizoğlu (2009): “Convergence in the finite Cournot oligopoly with social and individual learning,” *Journal of Economic Behavior and Organization*, 72, 670–690.

Vega-Redondo, F. (1997): “The Evolution of Walrasian Behavior,” *Econometrica*, 65, 375–384.

Vriend, N. J. (2000): “An illustration of the essential difference between individual and social learning, and its consequences for computational analyses,” *Journal of Economic Dynamics and Control*, 24, 1–19.

A Measure of Model Fit to the Experimental Data

The model's fit to the experimental data is measured based on the distribution of group-level total outputs. The unit of observations is the level of aggregate output in one period ($Q_t = \sum_i q_t^i$). Let $F_{sim,n}$ be the empirical distribution function based on the n observations generated by the simulation, and $F_{expr,m}$ be the empirical distribution function based on the m observations from the experiment.

$$F_{sim,n}(x) = \frac{1}{n}\sum_{i=1}^{n} I_{[-\infty,x]}(X_i)$$
$$F_{expr,m}(x) = \frac{1}{m}\sum_{i=1}^{m} I_{[-\infty,x]}(X_i)$$

where $I_{[-\infty,x]}(X_i)$ is the indicator function, equal to 1 if $X_i \leq x$ and equal to 0 otherwise, with $X_i \in \{100, 105, 110, \cdots, 395, 400\}$. Distance between the two distributions is measured by

$$D_{n,m} = \sum_x (F_{sim,n}(x) - F_{expr,m}(x))^6.$$

We also considered two other measures of distance

$$D2_{n,m} = \sum_x (F_{sim,n}(x) - F_{expr,m}(x))^2$$
$$D3_{n,m} = sup_x |F_{sim,n}(x) - F_{expr,m}(x)|$$

but the results are similar. These results are available from the authors upon request. Note that $D3_{n,m}$ places all the weight on minimizing the largest difference between the two distributions while sacrificing the distance in other parts. $D_{n,m}$ and $D2_{n,m}$, in contrast, take distances for all the bins into account.

B Two-step Grid Search

B.1 Coarse grid search

Grid searching is done for all possible combinations of shares of the five types in the population (f_I, f_C, f_R, f_F, and f_M for Individualistic, Cooperative, Retaliatory, Flexible, and Mixed, respectively) such that $f_x \in \{0, 0.2, 0.4, 0.6, 0.8, 1\}$ while $\sum_x f_x = 1$ where $x \in I, C, R, F, M$. For the other parameters, we consider the space of $B \in \{0, 20, 40, 60\}$. $G \in \{0, 20, 40, 60\}$, $\rho \in \{0.1, 0.3, 0.7, 0.9\}$, and $\omega^f \in \{0.1, 0.3, 0.7, 1\}$. The same exercise is carried out across treatments.

B.2 Finer grid search

B.2.1 Treatment Qq

The coarse grid search led to restricting our attention to the following subspace. For the shares of the five types $f_I \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$, $f_C \in \{0, 0.1, 0.2, 0.3, 0.4\}$, $f_R \in \{0, 0.1, 0.2, 0.3\}$ $f_F \in \{0, 0.1, 0.2, 0.3\}$, and $f_M = 1-(f_I+f_C+f_R+f_F)$. For the other parameters, $B \in \{30, 40, 50, 60, 80\}$, $G \in \{0, 20, 40, 60\}$, $\rho \in \{0.1, 0.3, 0.7, 0.9\}$ and $\omega^f \in \{0.1, 0.3, 0.7, 1\}$.

As a result, we obtain the following two sets of best parameter values.

	f_I	f_C	f_R	f_F	f_M	B	G	ρ	ω^f
Qq (a)	0.8	0.2	0.0	0.0	0.0	80	0	0	0
Qq (b)	0.7	0.3	0.0	0.0	0.0	30	0	0	0

B.2.2 Treatment $Qq\pi$

The coarse grid search led to restricting our attention to the following subspace. For the shares of the five types $f_I \in \{0.0, 0.1, 0.2, 0.3, 0.4\}$, $f_C \in \{0.0, 0.2, 0.3, 0.4, 0.5\}$,

$f_R \in \{0.0, 0.3, 0.4, 0.5, 0.6\}$ $f_F \in \{0.0, 0.2, 0.3, 0.4, 0.5, 0.6\}$, and $f_M = 1 - (f_I + f_C + f_R + f_F)$. For the other parameters, $B \in \{20, 30, 40, 50, 60, 80\}$, $G \in \{20, 30, 40, 50, 60\}$, $\rho \in \{0.1, 0.3, 0.7, 0.9\}$ and $\omega^f \in \{0.1, 0.3, 0.7, 1\}$.

As a result, we obtain the following set of best parameter values.

	f_I	f_C	f_R	f_F	f_M	B	G	ρ	ω^f
$Qq\pi$	0.0	0.2	0.4	0.4	0.0	60	50	0.1	1.0

C Supplementary Materials

Table 6: Effects of decision types on individual output.

Dependent variable: Average Individual Output			
	(1)	(2)	(3)
	Partial-Info	Full-Info	Full Sample
Cooperative	-.292***	-.26***	-.292***
	(.0298)	(.0378)	(.0284)
Retaliatory	.231***	.234***	.231***
	(.0244)	(.0349)	(.0233)
Flexible	.223	-.261	.223
	(.171)	(.203)	(.163)
Cooperative × FID			.0323
			(.0487)
Retaliatory × FID			.00321
			(.0433)
Flexible × FID			-.484*
			(.268)
FID			-9.59**
			(3.63)
Constant	70.8***	99.8***	70.8***
	(1.32)	(1.55)	(1.26)
Group FE	Yes	Yes	Yes
Observations	54	63	117

standard deviations in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Column 1 shows the statistics from the Partial-Info treatments. Column 2 shows the statistics from the Partial-Info treatments. Column 3 shows the statistics from the Full sample, where FID is a dummy variable: 1 if it is under Full-Info treatments, 0 otherwise.

D Experiment Introduction

D.1 Raven test

Instructions

Today you will be participated in the experiment on decision making. The experiment is anonymous, the data from your choices will only be linked to your station ID, not to your name. You will be paid for your participation. There is a show-up fee of $7.

The experiment consists of two parts. The first part involves an individual task that is described below. The second part is group decision making.

The individual task consists of 16 questions. You will be paid 25 cents for each question that you answer correctly. This amount will be added to your show up fee, and the earnings that you make in the second part of the experiment.

Your task is to analyze a geometric pattern and identify the missing part to complete the series. The pattern is in the form of 3x3 grid. You will have 10 minutes to answer 16 questions.

Next

We do not display the decision screen for the Raven test because of issues with copyright.

D.2 Cournot introduction

Here we show the on-screen Introduction for the 2-agent Full-Info treatment as an example. The players in the Partial-Info treatments read the same introduction but do not see anything about historical information.

Group Decision Making

Instructions

In this part of the experiment, the additional amount of cash that you earn will depend upon your decisions and the decisions of other participants. In this part, you will be earning experimental points. At the end of this part of experiment, the experimental points that you earned will be converted into dollars at the exchange rate of 5,000 experimental points = $1.

As your earnings depend on the decisions that you will make during the experiment, it is important to understand the instructions. Read them carefully. If you have any questions, raise your hand and the experimenter will come to your desk and provide answers.

You will make decisions for a firm in this experiment. You will be asked repeatedly to determine the quantity that your firm will produce. Your payoff will depend on your production and the production of 1 other firm(s). The decisions for these 1 other firm(s) will be made by 1 other participant(s) in the experiment.

You will be matched with the same 1 other firm(s) throughout the experiment. The same participants will make the decisions for these 1 firm(s). You will not know with whom you will be matched, like others will not know with whom they will be matched. Anonymity is ensured.

Own production and costs

The experiment will last for 100 periods. Each period you will decide how much your firm produces. Your production must be greater than or equal to 40 units and smaller than or equal to 170 units. You may only choose integer numbers. For example, it is not allowed to produce half units.

Producing involves costs. The cost function is

Your total cost = (Units produced by your firm)$^{1.5}$

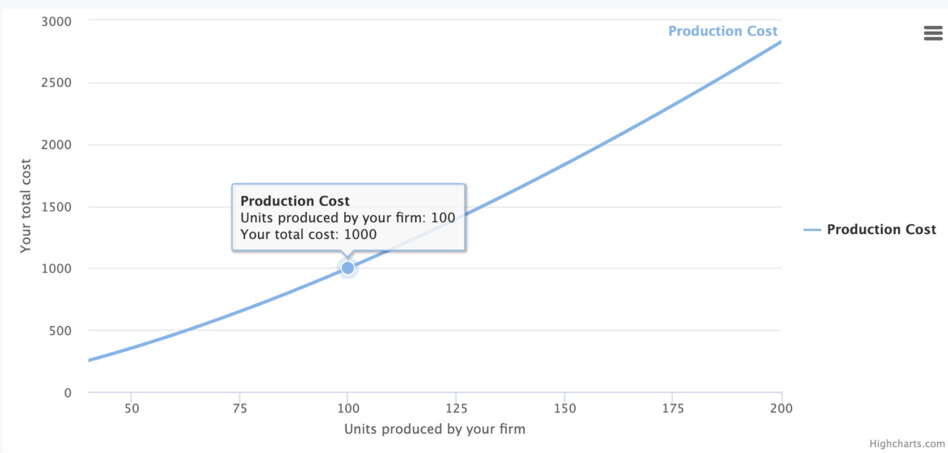

Question:

What will be your costs if you produce 100 units?

Your answer:

Total production and price

The price that you will receive for each unit produced depends on the total production. The total production is the sum of the production of all 2 firms. Since the other 1 firm(s) will also produce between 40 and 170 units, total production will not be smaller than 80 units and not be greater than 340 units.

The demand function is

Price = 45 - (3 * total production)$^{0.5}$

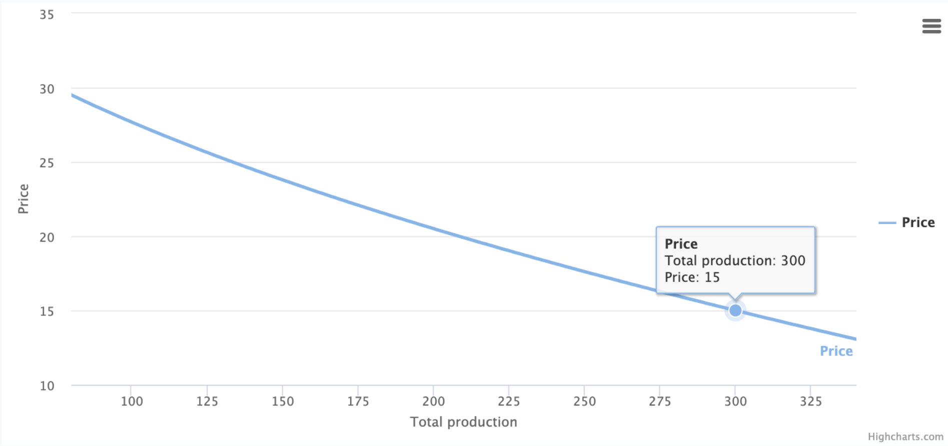

Question:

What will be the price if you produce 162 units, the other firm(s) produce 138 units?

Your answer:

Production and profit

Each period your revenue is equal to your produced quantity multiplied by the price. Your profit is your revenue minus the costs of your production.

The circumstances of production are exactly the same for other firm(s) as they are for you. The costs of their production is determined in a similar way as your production. They receive the same price for each unit they produce as you receive.

Profit Calculator

A profit calculator is provided. It can calculate the hypothetical profit, given the output of your firm and the other firm(s).

The experiment lasts for 100 periods. Each period the circumstances of production will be the same as described above. In this part of the experiment, your payoff is equal to the sum of your profits in each of the 100 periods. These profits are denoted in points.

Historical Information

When you are making decisions, you can click the **Show History** button to see the output in the previous periods.

At the end of this part of experiment, your points will be exchanged for real money. For each 5000 points you will receive one dollar, or 0.2 dollar per 1000 points.

Your total payoff in the experiment is equal to show up fee + earnings from the individual test + earnings from the group decision making.

D.3 Cournot decision page

Once players answer the question correctly, they will see the decision page, where they can move the slide bar to calculate the hypothetical payoff. The same instructions are available at the bottom.

Production

How many units will you produce (from 40 to 170)?

70

Next

Calculator for profit (move the slider to see your profit)

Total output of 1 other firm(s):

Output of Your firm:

40 170

Your profit is:

In the Full-Info treatments (but not the Partial-Info treatments), they can also click the "Show History" button to see the outputs of the other players.

Click the button to see the information in previous rounds: Show History

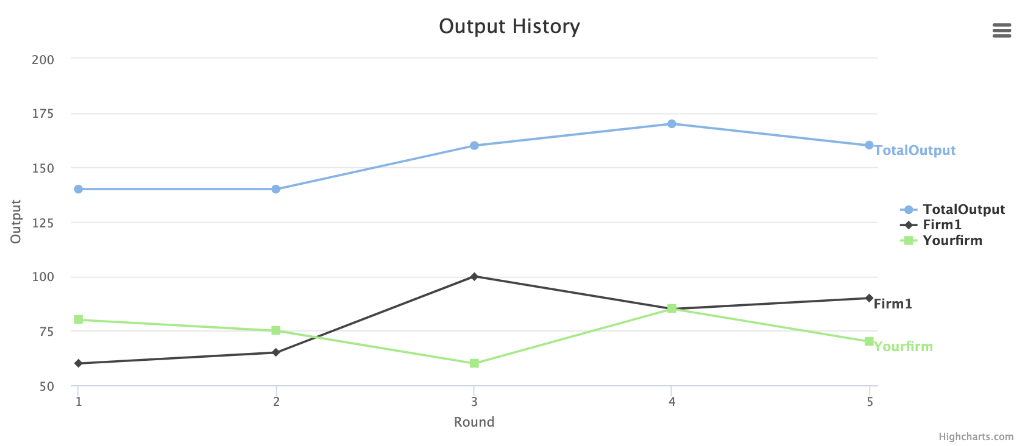

Round history

Round	Total Output	Price	Your Output	Your Profit	Firm1 Output	Firm1 Profit
5	160	23.09	70	1030.64	90	1224.29

D.4 Cournot result page

Once everyone in the group has made their decisions, they see the result page. The players in the Partial-Info treatment do not see the information of other firms.

Results

The results of round 5 are shown in the following table.

Your firm produced:	70 units
Total production:	160 units
Unit selling price:	23.09
Your profit:	1030.64
Your accumulated profit :	5507.31

Information of other firms

Firm 1 produced:	90 units
Firm 1 profits:	1224.29

Next