Available online at www.sciencedirect.com

ScienceDirect

ELSEVIER

Journal of Economic Theory 145 (2010) 1603–1617

JOURNAL OF Economic Theory

www.elsevier.com/locate/jet

Imitation and the evolution of Walrasian behavior: Theoretically fragile but behaviorally robust ☆

Jose Apesteguia [a], Steffen Huck [b], Jörg Oechssler [c,*], Simon Weidenholzer [d]

[a] *Department of Economics, Universitat Pompeu Fabra, Spain*
[b] *Department of Economics and ELSE, University College London, United Kingdom*
[c] *Department of Economics, University of Heidelberg, Germany*
[d] *Department of Economics, University of Vienna, Austria*

Received 23 April 2009; final version received 18 January 2010; accepted 20 February 2010

Available online 28 April 2010

Abstract

A well-known result by Vega-Redondo (1997) [18] implies that in symmetric Cournot oligopolies, imitation leads to the Walrasian outcome. We show that this result is not robust to the slightest asymmetry in costs, since *every* outcome where agents choose identical actions will be played some fraction of the time in the long run. We then conduct experiments to check this fragility. We obtain that, contrary to the theoretical prediction, the Walrasian outcome is a good predictor of market outcomes. Finally, we suggest a new theory based on a mix of imitation and other learning processes that explains subjects' behavior fairly well.

JEL classification: C72; C91; C92; D43; L13

Keywords: Evolutionary game theory; Stochastic stability; Imitation; Cournot markets; Information; Experiments; Simulations

☆ We are grateful to Klaus Abbink, Antonio Cabrales, Pedro Rey-Biel, Karl Schlag, Joep Sonnemans, and to audiences at the ESA Conference 2008 in Pasadena, the EEA-ESEM 2008 in Milan, and at the universities of Osnabrück and Vienna for very helpful comments and to Jon Benito for help with the simulations. We are also thankful to two anonymous referees and an associate editor for their very constructive feedback. Apesteguia acknowledges financial support by the Spanish Commission of Science and Technology (ECO2009-12836), the Barcelona GSE research network, and the Government of Catalonia. Weidenholzer acknowledges financial support from the Vienna Science and Technology Fund (WWTF) under project fund MA 09-017.

* Corresponding author. Fax: +49 6221 54 2997.

E-mail addresses: jose.apesteguia@upf.edu (J. Apesteguia), s.huck@ucl.ac.uk (S. Huck), oechssler@uni-hd.de (J. Oechssler), simon.weidenholzer@univie.ac.at (S. Weidenholzer).

0022-0531/$ – see front matter
doi:10.1016/j.jet.2010.02.014

1. Introduction

In an important paper, Vega-Redondo [18] shows how imitation of successful behavior can push agents towards very competitive outcomes. Specifically, he shows that in Cournot games imitation of the most successful strategies leads in the long run to the Walrasian outcome where price is equal to marginal cost.[1] This result is important since Cournot games serve as the main workhorse model for industrial organization. Furthermore, Cournot games reflect, more generally, environments where there is a tension between cooperation and competition, with the Cournot–Nash equilibrium outcome somewhere in between perfect collusion and perfect competition.

Two experimental papers (Huck, Normann, and Oechssler [7] and Offerman, Potters, and Sonnemans [12]) confirm the behavioral relevance of Vega-Redondo's [18] findings. When experimental subjects have access to information that allows them to imitate their rivals, competition gets significantly more intense. This is true even when subjects have all the necessary information to play the Nash equilibrium. In fact, both papers show that if subjects have just the necessary information to play a best reply, they converge to Cournot–Nash. However, additional information about rivals' choices and performance—which orthodox game theory deems irrelevant—leads them away from equilibrium play towards more competitive outcomes.[2]

In this paper we re-examine both, Vega-Redondo's [18] theoretical result and the experimental findings on it. First, we show that Vega-Redondo's theoretical result is surprisingly fragile. Slightest differences in costs are shown to have a huge impact on the long-run behavior of firms.[3] Specifically, we show that for an arbitrarily small change in some firm's cost, *every* outcome where firms choose identical actions will be played some fraction of the time in the long run if the action grid is not too coarse. The intuition for this result is simple. If a firm with a slight cost advantage moves to a slightly different quantity, it will, due to its cost advantage, still be the most successful firm and will thus be copied by others.

One might argue that imitation is a much less appealing concept in asymmetric settings. On a behavioral level, however, perfect symmetry is rare outside a model world where it is assumed. Thus, adaptive rules of behavior that hinge on perfect symmetry would probably be doomed to irrelevance when it comes to actual behavior by human subjects. Yet, imitation is well documented also outside the laboratory (see e.g. Miller and Dollard [10]).

We also show that this theoretical result is not only a curiosity that occurs in the limit. Rather we find in a series of simulations that small differences in agents' costs have large effects on all firms' behavior and profits if they imitate most of the time but experiment with some frequency. Specifically, we report that when one firm has a slight cost advantage, industry profits rise by more than 35% for experimentation rates of 10% or 20%.

Second, we conduct new experiments to analyze whether such cost differentials also change market outcomes in the laboratory. Our findings are clear-cut. Despite implementing a non-trivial

[1] See Alos-Ferrer and Ania [2] for a generalization of this result to a broader class of games.

[2] Since then this link between information, imitation, and competition has been replicated in a number of papers. See, for example, Abbink and Brandts [1], Huck, Normann, and Oechssler [8], or Selten and Apesteguia [15]. See also Apesteguia, Huck, and Oechssler [4] who analyze, both, theoretically and experimentally, the differences between Vega-Redondo's [18] model of imitation and Schlag's [13].

[3] Tanaka [17] studies imitation among firms with different (increasing) marginal costs. However, since in his setting, firms imitate only firms with the same cost structure, the long-run outcome is still such that price equals marginal costs for each firm.

cost differential, we find *no* change in outcomes. When subjects can observe their rivals, outcomes are far more competitive than predicted by the Cournot–Nash equilibrium, *regardless* of whether there are differences in costs or not. This confirms the strong behavioral link between feedback about rivals ("market transparency") and competitive behavior.

Finally, we suggest a new theory based on a mix of imitation and a number of other learning processes that is able to capture subjects' behavior. In particular, we show that a mix of imitation, best reply behavior, fictitious play, and relative payoff maximization matches our experimental data fairly well.

2. Theoretical predictions

As in Vega-Redondo's [18] model, we consider a market for a homogeneous good where a set of firms $N = \{1, \dots, n\}$ is competing à la Cournot. Each firm i produces some quantity q_i. The vector of quantities by firms other than i is denoted by q_{-i}. The total quantity $Q = \sum_{i=1}^{n} q_i$ produced by all firms determines the market price via a decreasing inverse demand function $0 \leqslant p(Q) < \infty$. In line with the prior literature, we assume that firms choose their output from a common grid $\Gamma = \{\delta, 2\delta, \dots, v\delta\}$ with $\delta > 0$ and $v \in \mathbb{N}$. We assume that v is large enough such that eventually $p(nv\delta) = 0$. We assume further that firms have cost functions that are weakly increasing and continuous for $q > 0$. Finally, we assume that firms 1 through $n-1$ have the same cost function $C_i(\cdot) = C(\cdot)$.

We consider two cases. In the symmetric case, firm n has the same cost function $C(\cdot)$. In the asymmetric case, firm n has a cost advantage, that is, $C_n(q) < C(q)$, $\forall q > 0$ and $C_n(0) \leqslant C(0)$. This formulation encompasses a possible difference in fixed costs as well as differences in marginal costs. Profits of firm i are given by

$$\pi_i(q_i, q_{-i}) = p(Q)q_i - C_i(q_i).$$

For the perfectly symmetric case in which all firms have the same costs, the (symmetric) Walrasian quantity q^w is defined as the quantity that every firm will produce in a market clearing situation when it takes the price as given. Formally, q^w is implicitly defined by

$$p(nq^w)q^w - C(q^w) \geqslant p(nq^w)q - C(q), \quad \forall q \geqslant 0.$$

We assume that there is a unique q^w,[4] and that $q^w \in \Gamma$, i.e. the Walrasian quantity is on the quantity grid. Note that when the cost function is differentiable and convex, then the Walrasian quantity is the quantity at which price equals marginal cost when all firms produce the same quantity, $p(nq^w) = C'(q^w)$.

After each period $t = 1, 2, \dots$ each firm observes the quantities produced and the profits associated with these quantities of all firms in the market. It then chooses the quantity that yielded the highest profit in the previous period. That is, we are considering an *imitate-the-best-max rule*.[5] More formally, in period t firm i chooses

$$q_i^t = q_j^{t-1} \quad \text{with } j \in \arg\max_{m \in N} \pi_m^{t-1}\left(q_m^{t-1}, q_{-m}^{t-1}\right).$$

Ties are assumed to be broken randomly. In addition, with a small probability $\varepsilon > 0$ each firm ignores the action prescribed by the imitation rule and chooses an action at random from all actions in Γ.

4 Standard assumptions like non-decreasing marginal costs and small fixed cost would guarantee this.

5 See Apesteguia et al. [4] for a discussion of various imitation rules.

The adjustment process described above gives rise to a Markov process. We use methods developed by Freidlin and Wentzell [6] (first applied in an economic context by Kandori, Mailath, and Rob [9]; Nöldeke and Samuelson [11]; and Young [19]) to identify the set of stochastically stable states, i.e. states that are in the support of the limit invariant distribution as the mutation probability ε goes to zero. We denote the monomorphic state in which all players set the same quantity q by ω_q and we shall call the state ω_{q^w} the Walrasian state.

Note that if all firms have identical cost functions and are in some absorbing state ω_q other than the Walrasian state, a single mutation towards the Walrasian quantity q^w is always imitated by other firms. The simple reason for this is that if price exceeds marginal cost, the firm with the highest quantity makes the largest profit and will be imitated. If prices are below marginal costs, the firm with the lowest quantity makes the smallest loss and will be imitated. Hence, as shown by Vega-Redondo [18], with identical cost functions only the state in which all firms set the Walrasian quantity is stochastically stable.

If however firm n has a cost advantage, it may be the case that after a mutation of firm n away from the Walrasian quantity it still earns the highest profit and hence will be imitated. Other firms, of course, do not realize that this higher profit is due to the lower cost. They simply observe that the strategy choice of firm n was more successful.

Proposition 1.

(1) *If there are no differences in costs, then the Walrasian state ω_{q^w} is the unique stochastically stable state.*
(2) *For any difference in costs and for any $q' > 0$, there exists a grid size δ^* such that for all $\delta < \delta^*$, all monomorphic states ω_q with $q \in \Gamma$, $q \geqslant q'$ are stochastically stable states.*
(3) *If there exists a difference in fixed costs such that $C(q) - C_n(q) \geqslant f$, $\forall q$, or if $C'(0) = \infty$, there exists a grid size δ^* such that for all $\delta < \delta^*$, the set of stochastically stable states is given by the set of all monomorphic states on the grid, $\{\omega_q \mid q \in \Gamma\}$.*

Proof. Part (1) follows without modification from Vega-Redondo [18].

With respect to Parts (2) and (3), note that as in Vega-Redondo's model, under the imitate-the-best-max rule only monomorphic states are absorbing. To see this point, consider any non-monomorphic state ω. Assume that firms make different profits and say firm j makes the highest profits. In the next round, all firms will imitate firm j and we reach the state ω_{q_j}. Note that there is also the (non-generic) case that firm j and firms $i \neq j$ make the same profits but offer different quantities. However, since ties are broken randomly, with positive probability the dynamics will shift us away from the non-monomorphic state.

Now, consider any monomorphic state ω_q and assume that firm n mutates to a new quantity $q + \delta$. Such a quantity increasing mutation will be followed if

$$p(nq + \delta)(q + \delta) - C_n(q + \delta) \geqslant p(nq + \delta)q - C(q)$$

or

$$\frac{C_n(q + \delta) - C(q)}{\delta} \leqslant p(nq + \delta). \tag{1}$$

Since for all $q \in \Gamma$, $C_n(q) - C(q) < 0$ by continuity of the cost function, we can always find a δ small enough such that the left-hand side is negative. Since the right-hand side is weakly positive, the inequality is satisfied for all q in the grid. Thus, for sufficiently small δ the process can always move *upward* with a chain of single mutations.

Next, consider any monomorphic state ω_q with $q \in \Gamma$, $q > \delta$. Assume that firm n mutates and decreases its quantity by the smallest possible unit, i.e. firm n mutates to $q - \delta$. This downward mutation will be followed if

$$p(nq-\delta)(q-\delta) - C_n(q-\delta) \geqslant p(nq-\delta)q - C(q)$$

or

$$\frac{C(q) - C_n(q-\delta)}{\delta} \geqslant p(nq-\delta). \tag{2}$$

As long as (2) is satisfied, the process can move downward with a chain of single mutations.

A sufficient condition for (2) to be satisfied is that

$$\lim_{\delta\to 0} \frac{C(q) - C_n(q-\delta)}{\delta} = \infty, \tag{3}$$

as $p(Q) < \infty$. For any fixed $q' > 0$, (3) holds for sufficiently small δ since $C(q') - C_n(q'-\delta) \geqslant C(q') - C_n(q') > 0$. Hence *downward* movements to the largest grid point less than or equal to q' are possible. Hence, for small enough δ the process can access all monomorphic states in the grid between q' and the upper bound of the grid from each other via a chain of single mutations. It follows that all these monomorphic states form one large "mutation connected component", which is stochastically stable (see Nöldeke and Samuelson [11]), which proves Part (2) of the proposition.

To prove Part (3), note that a first sufficient condition for (3) is that there is a difference in fixed costs of $C(q) - C_n(q) \geqslant f > 0$ between the cost functions since then

$$C(q) - C_n(q-\delta) \geqslant C(q) - C_n(q) \geqslant f > 0, \quad \forall q \in \Gamma.$$

A second sufficient condition for (3) is that $\lim_{q\to 0} C'(q) = \infty$. As explained above, for any fixed $q' > 0$ (3) holds for sufficiently small δ. Thus, it remains to show that (3) holds for $q \to 0$. Since

$$\lim_{\delta\to 0} \frac{C(q) - C_n(q-\delta)}{\delta} \geqslant \lim_{\delta\to 0} \frac{C(q) - C(q-\delta)}{\delta} = C'(q),$$

this is guaranteed by the assumption that $C'(0) = \infty$. Hence, for sufficiently small grid size δ the process can move to the lower bound of the grid with a chain of single mutations. Therefore, for small enough δ the process can access all monomorphic states in the grid from each other via a chain of single mutations, which proves Part (3). □

Hence, with asymmetric costs, the point prediction of Walrasian behavior breaks down.

3. Experimental design

In our experiment, subjects played repeated 3-player Cournot games in fixed groups for 60 periods. Marginal costs were set to 0. The only potential difference between firms were their fixed costs. The payoff function for each round was given by

$$\pi_i(q_i, q_{-i}) = p(Q)q_i - f_i,$$

with $p(Q) = \max\{120 - Q, 0\}$ being the inverse demand function.

The grid of quantities was given by $\Gamma = \{20, 21, 21.5, \ldots, 39.5, 40\}$.[6] Note that the symmetric joint profit maximizing output is at $q^c = 20$, the Cournot–Nash equilibrium output is at $q^N = 30$, and the symmetric Walrasian output is at $q^w = 40$.

In order to make imitation salient and give the theoretical results the best shot, subjects were not told anything about the game's payoff function apart from the fact that their payoff deterministically depended on their own choice and the choices of the two other subjects in their group, and that payoff functions and group compositions would not change over time.[7] After each period, subjects learned their own payoff, and the actions and payoffs of the two other subjects in their group. The 40 actions in Γ were labeled as $1, 2, \ldots, 40$ in ascending order.

We ran two treatments, one symmetric and one asymmetric, that differed only in the value of the f_i's. In treatment SYM, there were no fixed costs, $f_i = 0$ for all i. In treatment ASYM, however, there was a fixed bonus for firm 3, $g = -f_3 = 50$, while $f_i = 0$ for $i = 1, 2$. This amounts to the same as having a fixed cost of 50 for firms 1 and 2 but has the advantage of avoiding losses for subjects, which are difficult to enforce in an experiment. Subjects were not informed about differences in fixed cost in ASYM although they might have noticed them when all subjects in a group chose the same or similar actions but realized different payoffs.

The computerized experiments[8] were run in the ELSE laboratory at UCL. We had 7 independent groups in SYM and 8 in ASYM. In total 45 subjects participated in the experiment, drawn from the student population at UCL.[9] Subjects were paid a show-up fee of £5 and in addition to this were given £0.005 per point won during the experiment. The average payment was around £11 per subject, including the show-up fee. All sessions lasted less than 60 minutes.

Given this setup we can derive the following theoretical predictions from Proposition 1. In treatment SYM, the Walrasian quantity q^w is the unique stochastically stable state according to the imitate-the-best-max rule. However, in treatment ASYM, all monomorphic states ω_q with $20 < q \leqslant 40$ are in the support of the limit invariant distribution and should be observed with strictly positive probability in the long run.

To obtain quantitative predictions about profits in the short and medium run, we conducted computer simulations that allowed for various noise levels. The program followed with probability $1 - \varepsilon$ the imitation rule and chose actions with a uniform distribution from Γ with probability ε. In 10,000 repetitions of 60 periods, profits were 35.2% higher on average in ASYM than in SYM for $\varepsilon = 0.2$ and 37.8% higher for $\varepsilon = 0.1$. Notice that these differences were calculated excluding the bonus g paid to firm 3. This gives us a more realistic hypothesis for the experiment that does not simply compare the two long-run predictions.

Hypothesis P. Profits in treatment ASYM will be higher than in treatment SYM.

Fig. 1 shows the histograms for individual quantities in those simulations (with a noise level of 20%). Although 40 is the modal quantity in both settings, the distribution of quantities is much more spread out across the grid for ASYM, in particular between 30 and 40.

---

6 Quantity 20.5 was excluded to have exactly 40 strategies.

7 In order to test the robustness of the Walrasian outcome as the result of imitation, our experimental design aimed to give the best chances for imitation to play a significant role. Arguably, when subjects have no information about the payoff function, imitation is more salient.

8 The program was written with z-tree of Fischbacher [5].

9 We recruited 8 groups for both treatments but due to no-shows, only 7 groups were completed in SYM.

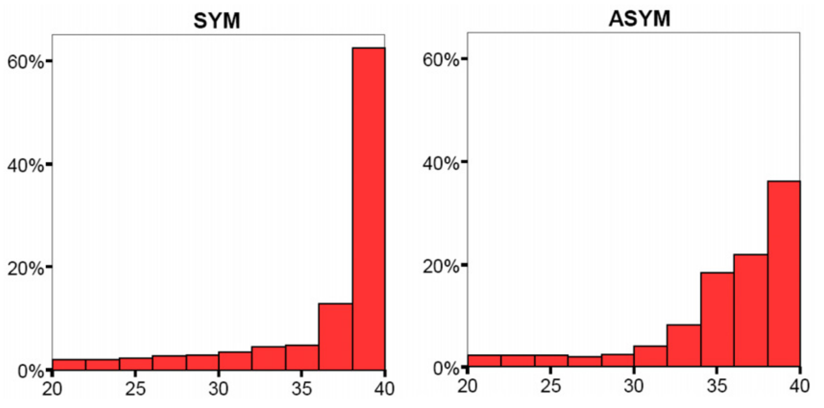

Fig. 1. Histograms of quantities in the simulations of the imitation process with a noise level of 20%. Note: Shown are results from 10,000 simulations of 60 rounds.

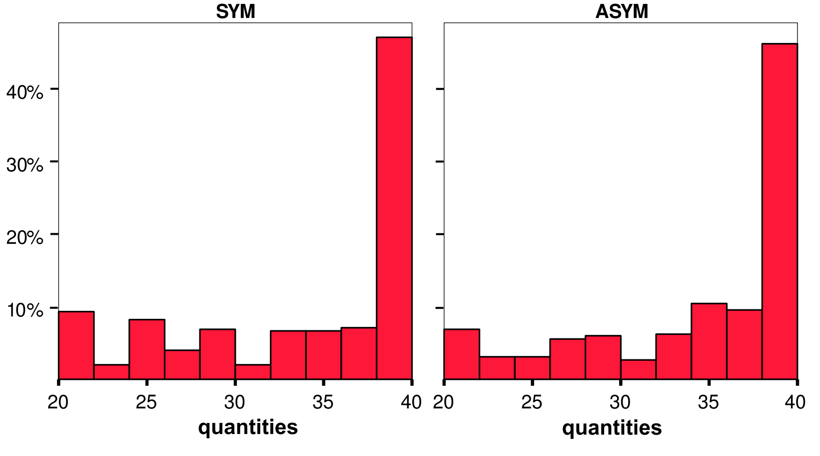

Fig. 2. Histograms of individual quantities by treatment.

4. Experimental results

Fig. 2 shows relative frequencies of actions separately for our two treatments. The two histograms look remarkably similar and indeed, there is no significant difference between the two distributions according to a Kolmogorov–Smirnov test (for a reference to all tests in this paper, see Siegel and Castellan [16]) even at the 10% significance level.[10]

Table 1 shows average quantities and the percentage deviation of average profits from the Cournot equilibrium profits for the two treatments over all periods.[11] Profits for treatment ASYM are calculated excluding the bonus of 50 for firm 3.

[10] From now on, whenever we say "not significant" we mean "not even significant at the 10% level."

[11] There is no noticeable time trend in the data.

Table 1
Summary statistics.

Treatment	Average quantities	% deviation from Cournot equilibrium profits
SYM	34.1	−39.8
ASYM	34.7	−42.2

Note: Profits in ASYM do not include g.

We find no significant difference between the distributions of quantities according to Mann–Whitney U tests on the basis of average quantities per group.[12] Likewise, there is no significant difference with respect to the deviation from Cournot profits. However, for both treatments we observe a sizable deviation from Cournot profits towards the zero-profit predictions of the competitive equilibrium. This seems remarkable given the understandable resistance of subjects to remain near this zero-profit area.

We summarize our results as follows.

Result. (a) Contrary to the theoretical prediction, there is no significant difference between our SYM and ASYM treatments in terms of quantities. In fact, in both treatments the mode of quantities is at the Walrasian quantity of 40.

(b) In both treatments there is a substantial deviation of profits from the Cournot equilibrium profit: Firms earn around 40% less. We find no support for Hypothesis P, which predicts higher profits in ASYM.

This leaves us with a puzzle to explain. Although data from treatment SYM are broadly consistent with theoretical predictions and simulations, data from treatment ASYM are not. To explain this, we would need a theory that predicts the same outcome in both treatments. In the next section we shall suggest one such theory that is based on a mix of several observed behavioral rules used by our subjects: imitation, myopic best reply, fictitious play, and relative payoff maximization.

According to a simple *myopic best reply* process, firms in period t best respond to the quantity set by the other firms in the previous period $t-1$. So given q_{-i}^{t-1}, in period t firm i chooses from

$$BR_i^t\left(q_{-i}^{t-1}\right) := \arg\max_{q_i \in \Gamma} \pi\left(q_i, q_{-i}^{t-1}\right).$$

Consequently in our setup the (myopic) best reply is given by[13]

$$BR_i^t\left(q_{-i}^t\right) = \frac{120 - \sum_{j \neq i} q_j^{t-1}}{2}.$$

Second, we consider a standard *fictitious play* process where firm i in period t best responds to the empirical distribution of total quantities of the other firms in the past $t-1$ periods. Given our linear Cournot setting, this is equivalent to best responding to the average quantities set by the other firms during the past $t-1$ periods. More formally, if we let $\hat{q}_j^{t-1} := \frac{1}{t-1}\sum_{k=1}^{t-1} q_j^k$ denote

[12] The cost advantaged players in treatment ASYM actually have a slightly lower average quantity (33.25) compared to the other firms (35.45) in this treatment. However, this difference is not significant at the 10% level of a MWU-test based on the 8 groups we have.

[13] If one of the learning processes predicts a quantity not on the grid, we round to the closest point on the grid.

the average quantity set by firm j during the past $t-1$ periods, then firm i chooses in period t from

$$FP_i^t\left(\hat{q}_{-i}^{t-1}\right) := \arg\max_{q_i \in \Gamma} \pi\left(q_i, \hat{q}_{-i}^{t-1}\right).$$

It follows in our setup that

$$FP_i^t(\hat{q}_{-i}) = \frac{120 - \sum_{j \neq i} \hat{q}_j^{t-1}}{2}.$$

Finally, we also consider *relative payoff maximization* where firms try to maximize relative payoffs, that is, the difference between their own payoff and the average payoff of the other firms given the other firms' quantities from the previous period,

$$RPM_i^t\left(q_{-i}^{t-1}\right) := \arg\max_{q_i \in \Gamma}\left[\pi_i\left(q_i, q_{-i}^{t-1}\right) - \frac{1}{n-1}\sum_{j \neq i} \pi_j\left(q_i, q_{-i}^{t-1}\right)\right]. \tag{4}$$

Note that in our setup the *relative best reply* to any quantity set by the other firms is the Walrasian quantity $q_w = 40$ in both treatments regardless of the other firms' quantities.[14]

We now assess the relevance of these theories with respect to individual behavior. Let K denote the set of the four theories considered so far. We shall do so by calculating for each round (except round 1) the quantity $q_i^t(k)$, which is predicted by theory $k \in K$ for player i in period t, and comparing it to the actually chosen quantity in that period, q_i^t. We can then calculate for each theory the squared deviation (SD),

$$SD_i^t(k) = \left(q_i^t(k) - q_i^t\right)^2.$$

We shall consider three different methods for aggregating the SD in order to classify subjects' behaviors. Method 1 calculates for each player and for each theory the mean squared deviation over all periods $t = 2, \ldots, 60$, and classifies a subject as best described by theory k if theory k minimizes the aggregated SD across all periods, i.e. if

$$k \in \arg\min_{k' \in K} \frac{1}{59}\sum_{t=2}^{60} SD_i^t\left(k'\right). \tag{MSD1}$$

We denote the percentage of subjects classified according to (MSD1) as $MSD1(k)$.

Method 2 counts for each player the number of periods in which a theory has the minimum SD and classifies subject i as best described by theory k if it does so for the largest number of periods,

$$k \in \arg\max_{k' \in K}\left|\left\{t\colon\, k' \in \arg\min_{k'' \in K} SD_i^t\left(k''\right)\right\}\right|. \tag{MSD2}$$

We denote the percentage of subjects classified according to (MSD2) as $MSD2(k)$.

[14] The relative profit of firm 1 is given by $RP_1 = (120 - q_1 - q_2 - q_3)q_1 - 0.5(120 - q_1 - q_2 - q_3)q_2 - 0.5(120 - q_1 - q_2 - q_3)q_3$. This function is strictly concave in q_1 and has a unique maximum at $q_1 = 60 - 0.25q_3 - 0.25q_2$. Since the strategy space is restricted to $\Gamma = \{20, 21, 21.5, \ldots, 39.5, 40\}$, the optimal action in this grid is 40 for all $q_2, q_3 \in \Gamma$. Obviously, adding or subtracting a constant to the profit functions in (4) does not change the maximizer. Hence fixed costs have no influence on the maximizer.

Table 2
Classifications of subjects' behavior in percent.

Measure	Theory			
	Imitation	Best reply	Fictitious	RPM
*MSD*1	.69 (.75)	.00 (.00)	.22 (.25)	.09 (.00)
*MSD*2	.77 (.75)	.02 (.00)	.11 (.25)	.10 (.00)
*MSD*3	.47 (.49)	.12 (.16)	.16 (.22)	.25 (.13)

Note: The values in parentheses refer to the cost advantaged players in ASYM. In case of ties according to one of the measures, we proceeded as follows: If there are m theories which are best for a given subject according to MSD1 or MSD2, each of these theories is assigned $1/m$ subjects. If there are m theories which are best for a given round according to MSD2 or MSD3, each of these theories is assigned $1/m$ rounds.

Finally, Method 3 calculates for each period t and individual i, which theory describes best the behavior of this subject in this period. The percentage of all cases in which theory k explains the behavior best is then referred to as $MSD3(k)$,

$$MSD3(k) = \frac{1}{59 * 45} \sum_i \left| \left\{ t \colon k \in \arg\min_{k' \in K} SD_i^t(k') \right\} \right|. \tag{MSD3}$$

Table 2 shows the resulting classifications. Notice that the classifications differ somewhat depending on the measure used. However, according to all measures, imitation is the dominant mode of behavior. Yet the other theories also explain part of subjects' behavior and as we shall see in the next section, this can have important implications for the theoretical properties of a mixed process. It is also interesting to note that the cost advantaged players do not seem to behave differently from the other players as shown by the percentages in parentheses in Table 2. We find no significant difference between the cost advantaged player and the other two players, neither on the values of MSD1 nor on those of MSD2 (Wilcoxon tests).[15] Likewise, a MWU-test finds no significant difference between the SYM and ASYM groups, neither for MSD1 nor for MSD2.

5. An alternative theory

The goal of this section is to come up with a theory that matches the experimental results. In particular, two features need to be explained, namely that (1) the behavior in treatments SYM and ASYM is roughly the same, and (2) that observed quantities are distributed over the entire grid with a mode at 40.

One possible candidate for explaining the first feature is a process purely based on relative payoff maximization. It is well known (see e.g. the survey by Alos-Ferrer and Schlag [3]) that in a symmetric Cournot oligopoly, the long-run outcome of imitation corresponds to a finite population ESS (Schaffer [14]) and the latter, in turn, is characterized by maximization of relative payoffs. So far it was difficult to disentangle imitation and relative payoff maximization (RPM) as most studies were conducted in symmetric environments. Interestingly, the correspondence between imitation and relative payoff maximization breaks down for asymmetric games. As already noted, the relative payoff maximizing action is in our setup always $q^w = 40$ regardless of the quantities set by the other firms. Hence, the unique symmetric Nash equilibrium in the game where relative payoffs are maximized (and, thus, the finite population ESS in the original game)

[15] The Wilcoxon test cannot be applied in the case of MSD3 since it gives the average across all subjects.

is given by the Walrasian quantity q^w regardless of the treatment. Thus, RPM could explain the fact that behavior was similar in SYM and ASYM. However, RPM cannot explain why firms choose quantities below 40. Furthermore, Table 2 shows that RPM is an important but not the dominant determinant of firms' behavior. Therefore, we disregard this explanation.

In fact, we have seen from Table 2 that all four considered learning processes—imitation, best reply, fictitious play, and relative payoff maximization—may play a role in explaining individual behavior. Consequently, we are interested in the theoretical properties of a mixed process where firms choose their quantities according to either of these behavioral models at a given point in time. In particular, we assume that in a given period a firm chooses its quantity according to imitation with probability $w_1 > 0$, according to a best reply process with probability $w_2 > 0$, according to fictitious play with probability $w_3 > 0$, and according to relative payoff maximization with probability $w_4 > 0$, with $\sum w_i \leqslant 1$. There may or there may not be inertia in firms' decisions, i.e. $\sum w_i = 1$ or $\sum w_i < 1$.

Proposition 2. *For both treatments, SYM and ASYM, the mixed process of imitation, best reply, fictitious play, and RPM does not have any absorbing state. Instead there is only one large absorbing set. In particular, every monomorphic state is contained in the absorbing set.*

Proof. Consider any state ω. With positive probability, firm i chooses its quantity according to RPM, i.e. it chooses a quantity of 40. With positive probability the other firms either also choose their quantity according to RPM or they imitate firm i and we reach the state $\omega_{q^w} = (40, 40, 40)$. Consequently, the Walrasian state can be reached from any other state and the process may (due to imitation, RPM, or inertia) spend an arbitrary number of periods there. Now, with positive probability all firms choose their quantity according to best reply and, since $BR_i^t(40, 40) = 20$ we reach the state $(20, 20, 20)$. Again, due to imitation or inertia, the process may spend an arbitrarily large number of periods in this state. We can now justify any $\hat{q}_i$ as a time average of our process as it may spend an arbitrary number of periods in both the state $(20, 20, 20)$ and the state $(40, 40, 40)$. Next note that $FP_i(40, 40) = 20$ and $FP_i(20, 20) = 40$ and that $FP_i(\hat{q}_{-i})$ is continuously decreasing in $\sum_{j \neq i} \hat{q}_i$. Consequently, a firm adjusting its quantity according to fictitious play may set any quantity in Γ. In other words, all monomorphic states will be visited with positive probability. □

Note that the argument presented above is independent of the treatment. That is why we obtain the same qualitative prediction for both treatments. The exact quantitative prediction of this process will however depend on the relative influence of each of the behavioral components: RPM pushes the process towards the Walrasian quantity, whereas fictitious play and best reply push the process away from the Walrasian quantity. Imitation, may temporarily stabilize the process at any monomorphic state.

In order to obtain quantitative predictions, we have simulated such a mixed process with different sets of weights on the respective behavioral models. We have run simulations with exactly the weights as suggested by the different MSD measures from Table 2. Since experimental behavior is always subject to unpredictable noise, we have also run simulations in which players follow the respective learning process with probability 0.8 and choose randomly a quantity according to a uniform distribution from Γ with probability 0.2.[16]

[16] We have also run simulations for a noise level of 10%. The results are always between those for the 0% and 20% noise levels.

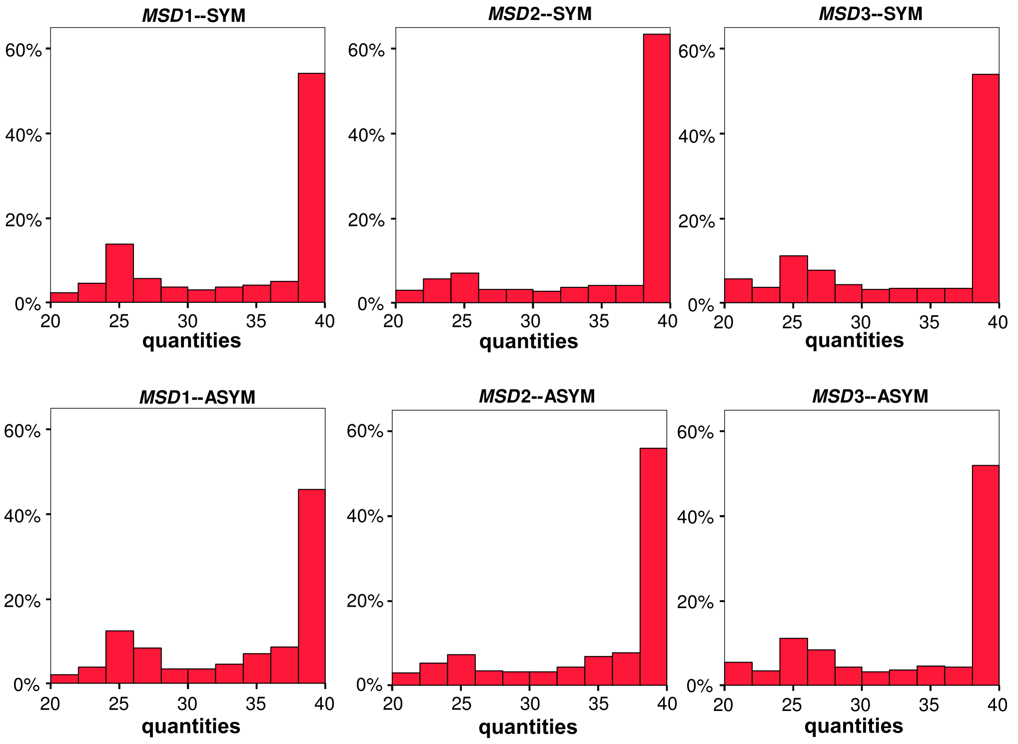

Fig. 3. Histograms of quantities in the simulations of the mixed process with weights according to the respective MSD measure and a noise level of 20%.
Note: Shown are results from 10,000 simulations of 60 rounds.

For the resulting distributions of quantities see the histograms in Figs. 3 and 4. The average quantities are presented in Table 3. The results seem to be fairly robust across the different weights used and the noise levels. The suggested mixed process tracks the experimental data remarkably well. In all cases, quantities are distributed over the entire strategy set and the histograms are very close to what we find in the experimental data. Furthermore, in all cases, the mode is at the Walrasian quantity of 40. In particular for the case with 20% noise, average quantities are very close to the ones found in the experiment. Profit levels in Table 4 are also fairly close to the ones observed in the experiment. Relative to the Cournot profit, deviations range from −40 to −44 percent (−50 to −57 percent, respectively) for MSD1 and MSD3 and a noise level of 20% (0%). Only for MSD2 are the deviations somewhat larger (up to −71).

It is also interesting to compare Figs. 3 and 4 with Fig. 2. Obviously, Fig. 3 is closer to Fig. 2 than is Fig. 4. Thus, with noise the mixed process predicts behavior better than without. However, this does not imply that we consider all deviations from 40 just noise. Even without noise, about 25% of quantities are less than 40. This agrees with Proposition 2, where it is shown that every quantity in the grid can occur under the mixed process.

Furthermore, observe that the noise levels in Figs. 1 and 3 are the same, yet the outcome differs quite a bit. This shows that noise and imitation alone are not sufficient to explain well the shape of the histograms of the experiment. The mixed process does much better (with and

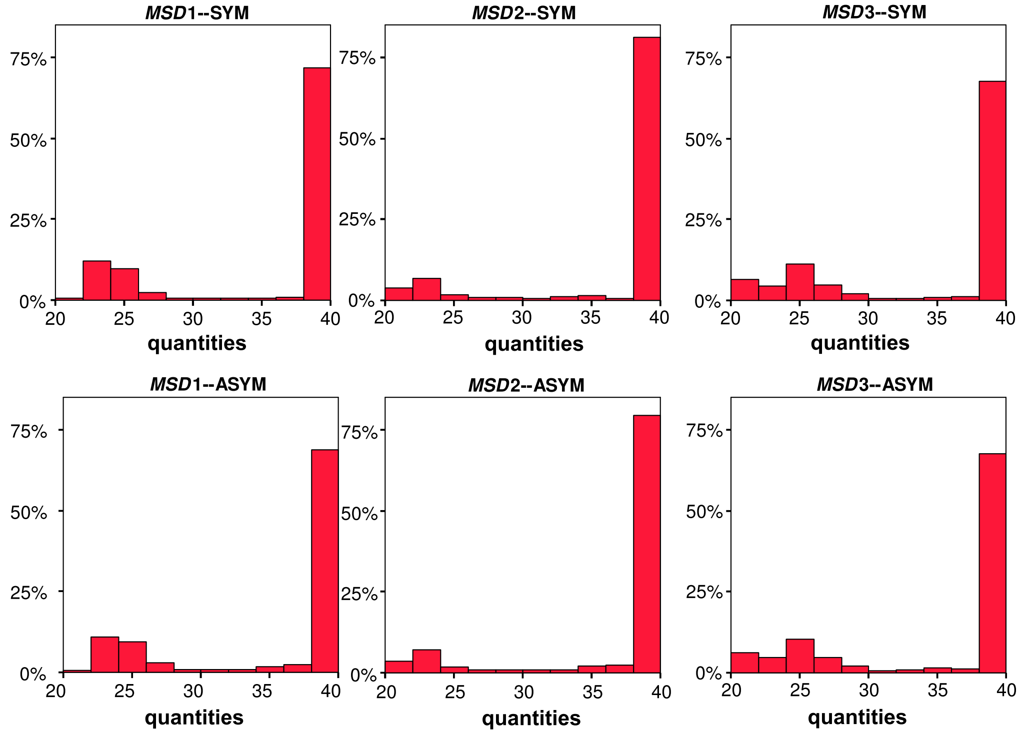

Fig. 4. Histograms of quantities in the simulations of the mixed process with weights according to the respective MSD measure and a noise level of 0%.
Note: Shown are results from 10,000 simulations of 60 rounds.

Table 3
Average quantities in the simulations.

Noise level	*MSD*1		*MSD*2		*MSD*3	
	SYM	ASYM	SYM	ASYM	SYM	ASYM
0%	35.90	35.85	37.35	37.30	35.29	35.39
20%	34.78	34.43	35.85	35.53	34.39	34.36

Table 4
Average percentage deviation from Cournot profits in the simulations.

Noise level	*MSD*1		*MSD*2		*MSD*3	
	SYM	ASYM	SYM	ASYM	SYM	ASYM
0%	−57.2	−56.4	−71.7	−70.8	−50.9	−51.7
20%	−44.9	−41.0	−55.2	−51.8	−41.3	−40.8

without noise) in particular with respect to two aspects: (1) the distributions for SYM and ASYM are similar and (2) there is a sizable fraction of choices at low quantities between 20 and 25. It seems that noise and the mixed process interact to produce this result.

6. Conclusion

In this paper we study the fragility and robustness of the prediction in Vega-Redondo's imitation theory [18]. If agents can observe their rivals and imitate the action that in the previous round was most successful, Walrasian outcomes emerge in the long run. However, as we show, in theory this does no longer hold if there are differences in costs, even if these differences are very small. Intuitively, one would think that such a fragility would severely limit the theory's predictive power. But intuition is wrong. Despite its theoretical fragility, the link between information about rivals and intense competition is robust. Differences in costs do not help subjects to overcome cut-throat competition. Somewhat ironically, the underlying reason for this stability of Vega-Redondo's prediction is that while imitation is the dominant mode of behavior when others' actions and profits are observable, subjects also use other adaptive rules. Fundamentally, our result stresses the behavioral importance of information about rivals that orthodox game theory deems irrelevant.

Appendix A. Instructions

Welcome to our experiment! Please read these instructions carefully. Do not talk with others and remain quiet during the entire experiment. If you have any questions, please ask us. We will come to you and answer your question privately.

During this experiment, which lasts for 60 rounds, you will be able to earn points in every round. You will form a group with two other participants. The composition of your group remains constant throughout the course of the experiment. The number of points you may earn depends on your action and the actions of the two other participants in your group. At the end of the experiment your accumulated points will be converted to pound sterling at a rate of 200 : 1.

Each round, you will have to choose one of 40 different actions, actions $1, 2, 3, \dots, 40$. Actions are ordered such that action 1 is the smallest and action 40 is the largest action. We are not going to tell you how your payoff is calculated, but in every round your payoff depends uniquely on your own decision and the decisions of the two other participants in your group. The rule underlying the calculation of the payoff does not depend on chance and remains the same in all 60 rounds.

After every round you get to know how many points you earned with your action in the current round. In addition, you will receive information about the actions of the other two participants in your group, and how many points each of them earned.

After the last period you will be reminded of all your 60 payoffs and the computer will calculate the sum of these which will then be converted into pound sterling.

These are all the rules. Should you have any questions, please ask now. Otherwise have fun in the next 60 rounds.

References

[1] K. Abbink, J. Brandts, 24. Pricing in Bertrand competition with increasing marginal costs, Games Econ. Behav. 63 (2008) 1–31.
[2] C. Alos-Ferrer, A. Ania, The evolutionary stability of perfectly competitive behavior, Econ. Theory 26 (2005) 497–516.
[3] C. Alos-Ferrer, K.H. Schlag, Imitation and learning, in: P. Anand, P. Pattanaik, C. Puppe (Eds.), The Handbook of Rational and Social Choice, Oxford University Press, 2009.
[4] J. Apesteguia, S. Huck, J. Oechssler, Imitation: Theory and experimental evidence, J. Econ. Theory 136 (2007) 217–235.

[5] U. Fischbacher, z-Tree: Zurich toolbox for ready-made economic experiments, Exper. Econ. 10 (2007) 171–178.
[6] M. Freidlin, A. Wentzell, Random Perturbations of Dynamical Systems, Springer-Verlag, Berlin, 1984.
[7] S. Huck, H.T. Normann, J. Oechssler, Learning in Cournot oligopoly: An experiment, Econ. J. 109 (1999) C80–C95.
[8] S. Huck, H.T. Normann, J. Oechssler, Does information about competitors' actions increase or decrease competition in experimental oligopoly markets? Int. J. Ind. Organ. 18 (2000) 39–57.
[9] M. Kandori, G. Mailath, R. Rob, Learning, mutation, and long run equilibria in games, Econometrica 61 (1993) 29–56.
[10] N.E. Miller, J. Dollard, Social Learning and Imitation, Yale University Press, New Haven, 1962.
[11] G. Nöldeke, L. Samuelson, An evolutionary analysis of backward and forward induction, Games Econ. Behav. 5 (1993) 425–454.
[12] T. Offerman, J. Potters, J. Sonnemans, Imitation and belief learning in an oligopoly experiment, Rev. Econ. Stud. 69 (2002) 973–997.
[13] K.H. Schlag, Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits, J. Econ. Theory 78 (1998) 130–156.
[14] M.E. Schaffer, Evolutionary stable strategies for a finite population and variable contest size, J. Theoret. Biol. 132 (1988) 469–478.
[15] R. Selten, J. Apesteguia, Experimentally observed imitation and cooperation in price competition on the circle, Games Econ. Behav. 51 (2005) 171–192.
[16] S. Siegel, J.Jr.N. Castellan, Nonparametric Statistics for the Behavioral Sciences, McGraw–Hill, Singapore, 1988.
[17] Y. Tanaka, Long run equilibria in an asymmetric oligopoly, Econ. Theory 14 (1999) 705–715.
[18] F. Vega-Redondo, The evolution of Walrasian behavior, Econometrica 65 (1997) 375–384.
[19] H.P. Young, The evolution of conventions, Econometrica 61 (1993) 57–84.