

University of Heidelberg

Department of Economics



Discussion Paper Series | No. 461

**Imitation and the Evolution of Walrasian
Behavior: Theoretically Fragile but
Behaviorally Robust**

Jose Apesteguia,
Steffen Huck,
Jörg Oechssler, and
Simon Weidenholzer

November 2007

Imitation and the Evolution of Walrasian Behavior: Theoretically Fragile but Behaviorally Robust

Jose Apestegua* Steffen Huck† Jörg Oechssler‡
Simon Weidenholzer§

November 30, 2007

Abstract

A well-known result by Vega-Redondo implies that in symmetric Cournot oligopoly, imitation leads to the Walrasian outcome where price equals marginal cost. In this paper we show that this result is not robust to the slightest asymmetry in fixed costs. Instead of obtaining the Walrasian outcome as unique prediction, *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 still a good predictor of behavior.

JEL codes: C72; C91; C92; D43; L13.

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

*Department of Economics, Universitat Pompeu Fabra. jose.apestegua@upf.edu

†Department of Economics and ELSE, University College London. s.huck@ucl.ac.uk

‡Department of Economics, University of Heidelberg. oechssler@uni-hd.de

§Department of Economics, University of Vienna. simon.weidenholzer@univie.ac.at

1 Introduction

In a seminal paper Vega-Redondo (1997) 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. This result is important since Cournot games not only serve as the main workhorse model for industrial organization but 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.¹

Two experimental papers (Huck, Normann, and Oechssler, 1999 and Offerman, Potters, and Sonnemans, 2002) confirm the behavioral relevance of Vega-Redondo’s 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 while subjects converge to Cournot-Nash if they have just the necessary information to play a best reply, 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 theory 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 agents. Specifically, we show that for an arbitrarily small change in some agent’s fixed costs, we can find an action set from which quantities are chosen³ such that *every* outcome

¹See Alos-Ferrer and Ania (2005) for a generalization of the result to a broader class of games.

²Since then this link between information, imitation and competition has been replicated in a number of papers. See, for example, Abbink and Brandts (2007), Huck, Normann, and Oechssler (2000), or Selten and Apesteguia (2005). See also Apesteguia, Huck, and Oechssler (2007) who analyse, both, theoretically and experimentally, the differences between Vega-Redondo’s (1997) model of imitation and Schlag’s (1998).

³More technically speaking, we can find a grid that is fine enough. Almost all models of imitation deal with finite action spaces.

where agents choose identical actions will be played some fraction of the time in the long run. The intuition for this is simple. If a firm with a slight fixed 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.

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 their profits if they imitate most of the time but experiment with a reasonable 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 behavior of subjects in the laboratory. Our findings are very clear-cut. Despite implementing a non-trivial 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.

2 Theoretical predictions

As in Vega-Redondo's (1997) model we consider a market for a homogeneous good where a set of firms $N = \{1, \dots, n\}$ is competing á la Cournot. Each firm i produces some quantity q_i . The vector of quantities by firms other than i is denoted by q_{-i} . In line with the prior literature, we assume for technical reasons that firms choose their output from a common grid $\Gamma = \{0, \delta, 2\delta, \dots, \nu\delta\}$ with $\delta > 0$ and $\nu \in \mathbb{N}$. The total quantity $Q = \sum_{i=1}^n q_i$ produced by all firms determines the price on the market via a linear inverse demand function $p(Q) = \max\{a - bQ, 0\}$. We assume that all firms face constant marginal costs c with $0 \leq c < a$. In addition, we assume that firm i may have to bear some fixed cost (or bonus) f_i . So firm i 's cost function is given by $c_i(q_i) = cq_i + f_i$. The fixed costs may differ among firms. Profits

of firm i are given by

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

The (symmetric) Walrasian quantity q^w is defined as the quantity at which price equals marginal cost c when all firms produce the same quantity. Within our setup we have

$$q^w = \frac{a - c}{bn}.$$

We assume that $q^w \in \Gamma$, i.e. the Walrasian quantity is contained in the quantity grid.

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*.⁴ More formally in period t firm i chooses

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

Ties are assumed to be broken randomly. In addition, with 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 . Let ω_q denote the monomorphic state where all players set the same quantity q .

The adjustment process described above gives rise to a Markov process. We use methods developed by Freidlin and Wentzel (1984) (first applied in an economic context by Kandori, Mailath, and Rob, 1993; Nöldeke and Samuelson, 1993; and Young, 1993) 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.

Let us now assume that some firm k has a cost advantage over all other firms in the market. We model this cost advantage via the fixed cost. In particular and without loss of generality, we assume that $f_i = 0$ for all $i \neq k$ and $-f_k = g \geq 0$.

⁴See Apesteguia et al. (2007) for a discussion of various imitation rules.

Note that if $g = 0$, i.e. all firms have identical cost functions, a single mutation towards the Walrasian quantity q^w is always imitated by other firms. The simple reason for this is that if the 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 largest profit and hence will be imitated. Hence, as shown by Vega–Redondo (1997), with identical cost functions only the state where all firms set the Walrasian quantity is stochastically stable.

If however firm k has a cost advantage, it may be the case that after a mutation of firm k 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 fixed cost. They simply observe that the strategy choice of firm k was more successful. This introduces another source of bounded rationality which pushes the system away from the Walrasian quantity.

Proposition 1 (1) *If there are no differences in fixed cost ($g = 0$), then the Walrasian state ω_{q^w} is the unique stochastically stable state.*

(2) *For any difference in fixed costs $g > 0$, 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, $\{\omega_q | q \in \Gamma\}$.*

Proof. The first part follows without modification from Vega–Redondo (1997).

With respect to the second part, note that as in Vega–Redondo’s model, under the imitate the best rule only monomorphic states are absorbing. Consider any non-monomorphic state ω . Assume that firms make different profits and say firm j makes the highest profits. With positive probability 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 k and firms $i \neq k$ make the same profits but offer different quantities. However, since ties are broken randomly, with positive probability the dynamics will shift us to the state ω_{q_j} .

We now identify the set of stochastically stable states for arbitrary g and δ . Consider some monomorphic state ω_q and assume that firm k mutates

and decreases its quantity by the smallest possible unit, i.e. firm k mutates to $q_k - \delta$. This (downward) mutation will be followed if firm k 's profits after the mutation exceeds the profits of the other firms, i.e. if and only if

$$((a - b(nq_k - \delta)) - c)(q_k - \delta) + g \geq ((a - b(nq_k - \delta)) - c)q_k.$$

So, a single downward mutation is followed if

$$q \geq q^w + \frac{\delta}{n} - \frac{g}{\delta bn} =: q^{low}. \quad (1)$$

Note that this implies that the lowest quantity that can be reached by a chain of single downward mutations is $q^{low} - \delta$. Obviously, from all $q > q^w$, a downward move is always possible, just like in Vega–Redondo (1997). But for $g > 0$, downward moves become possible for some $q < q^w$ as well.

Likewise, note that a single upward mutation $q_k + \delta$ of firm k is followed if

$$q \leq q^w - \frac{\delta}{n} + \frac{g}{\delta bn} =: q^{high}. \quad (2)$$

as long as $p > 0$. Again, we can move up to $q^{high} + \delta$ by a chain of single mutations.

Consider now the case $p = 0$, i.e. $q \geq \frac{a}{bn}$. An upward mutation is followed if $-c(q + \delta) + g \geq -cq$. That is if

$$\delta \leq \frac{g}{c}. \quad (3)$$

Note that if $q^{high} + \delta \geq \frac{a}{bn}$, inequality (3) holds also. That is, if we can move up to the point where the price is zero, we can move up all the way to the upper bound of our grid.

Figure 1 summarizes the results so far. All one-step mutations toward q^w are always possible. Downward movements for $q < q^w$ are possible if and only if (1) is satisfied. Upwards movements for $q^w < q < \frac{a}{bn}$ are possible if and only if (2) is satisfied. Upwards movements for $q > \frac{a}{bn}$ are possible if $q^w + \frac{\delta(n-1)}{n} + \frac{g}{\delta bn} \geq \frac{a}{bn}$ holds.

So all states in the following set can be reached from any other state by a series of single mutations

$$B = \left\{ \omega_q \mid q \in \Gamma, q^w - \frac{\delta(n-1)}{n} - \frac{g}{\delta bn} \leq q \leq \bar{q} \right\}$$




Figure 1: Transitions from one monomorphic state to a neighboring one that can be reached with one mutation.

where

$$\bar{q} = \begin{cases} v\delta & \text{if } q^w + \frac{\delta(n-1)}{n} + \frac{g}{\delta bn} \geq \frac{a}{bn} \\ q^w + \frac{\delta(n-1)}{n} + \frac{g}{\delta bn} & \text{else} \end{cases}$$

Hence, all states in B form one large “mutation connected component”, which is stochastically stable (see Nöldeke and Samuelson, 1993). Note that as $\delta \rightarrow 0$ the set B converges to the set $\{\omega_q | q \in \Gamma\}$. ■

3 Experimental design

In our experiment, subjects played repeated 3-player Cournot games in fixed groups for 60 periods. 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. Marginal cost were set to 0.

The grid of quantities was given by $\Gamma = \{20, 21, 21.5, \dots, 39.5, 40\}$.⁵ 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 $q^w = 40$.

⁵Quantity 20.5 was excluded to have exactly 40 strategies.

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 the payoff function was the same throughout all of the experiment. 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 $1, 2, \dots, 40$ in ascending order.

We ran two treatments, one symmetric and one asymmetric, that differed only on 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 is 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 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 are not informed about differences in fixed cost in ASYM although they may notice them when all subjects in a group choose the same or similar actions but realize different payoffs.

The computerized experiments⁶ 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.⁷ 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 hypothesis from Proposition 1.

Hypothesis Q 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 $\{\omega_q | q \in \Gamma\}$ are in the support of the limit invariant distribution and should be

⁶The program was written with z-tree of Fischbacher (2007).

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

observed with strictly positive probability in the long run.

To obtain quantitative predictions about profits in the short and medium run, we have conducted computer simulations that allow for realistic 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 $\varepsilon = 0.1$.⁸

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

4 Experimental results

Figure 2 shows relative frequencies of actions separately for our two treatments. There is clearly no significant difference between the two distributions according to a Kolmogorov–Smirnov test at any conventional significance level. The mode of both distributions is at 40, the Walrasian quantity, which is predicted by theory for SYM but not necessarily for ASYM.

Table 1 shows average quantities and the percentage deviation of average profits from the Cournot equilibrium profits for the two treatments over all periods.⁹ Profits for treatment ASYM are calculated excluding the bonus of 50 for firm 3.

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 .

⁸Profits in ASYM are calculated excluding the bonus of 50 for firm 3.

⁹There is no noticeable time trend in the data.




Figure 2: Histograms of individual quantities by treatment.

We find no significant difference between average quantities according to MWU tests (see, e.g., Siegel and Castellan 1988) on the basis of average quantities per group. 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 (1) 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 competitive quantity of 40.

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

5 Conclusion

In this paper we study the fragility and robustness of the prediction in Vega-Redondo's imitation theory. 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, 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. This stresses the behavioral importance of information about rivals that orthodox game theory deems irrelevant.

References

- [1] Abbink, K., and Brandts, J. (2007), "24", *Games and Economic Behavior*, forthcoming.
- [2] Alos-Ferrer, C. and Ania, A. (2005), "The Evolutionary Stability of Perfectly Competitive Behavior", *Economic Theory*, 26, 497–516.
- [3] Apesteguia, J., Huck, S. and Oechssler, J. (2007), "Imitation: Theory and Experimental Evidence", *Journal of Economic Theory*, 136, 217-235.
- [4] Fischbacher, U. (2007), "z-Tree: Zurich Toolbox for Ready-made Economic Experiments", *Experimental Economics*, 10, 171-178.
- [5] Huck, S., Normann, H.T., and Oechssler, J. (1999), "Learning in Cournot Oligopoly: An Experiment", *Economic Journal*, 109, C80-C95.
- [6] Huck, S., Normann, H.T., and Oechssler, J. (2000), "Does Information about Competitors' Actions Increase or Decrease Competition in

- Experimental Oligopoly Markets?”, *International Journal of Industrial Organization*, 18, 39-57.
- [7] Kandori, M., Mailath, G. and Rob, R. (1993), “Learning, Mutation, and Long Run Equilibria in Games”, *Econometrica*, 61, 29-56.
- [8] Offerman, T., Potters, J., and Sonnemans, J. (2002), “Imitation and Belief Learning in an Oligopoly Experiment”, *Review of Economic Studies*, 69, 973-997.
- [9] Schlag, K. (1998), “Why Imitate, and If So, How? A Boundedly Rational Approach to Multi-armed Bandits”, *Journal of Economic Theory* 78, 130-56.
- [10] Selten, R., and Apesteguia, J. (2005), “Experimentally Observed Imitation and Cooperation in Price Competition on the Circle”, *Games and Economic Behavior*, 51, 171-192.
- [11] Siegel, S. and N. Castellan, J. Jr. (1988), *Nonparametric Statistics for the Behavioral Sciences*, Singapore: McGraw-Hill.
- [12] Vega-Redondo, F. (1997), “The Evolution of Walrasian Behavior”, *Econometrica*, 65, 375-384.
- [13] Young, H.P. (1993), “The Evolution of Conventions”, *Econometrica*, 61, 57-84.

Appendix: 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, ..., 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.