

The “Win-Continue, Lose-Reverse” Rule in Cournot Oligopolies: Robustness of Collusive Outcomes

Segismundo S. Izquierdo and Luis R. Izquierdo

1 Introduction and Motivation

It is generally recognised that the actual decision-making processes followed by real-world firms when they have to set prices or production levels have often little to do with those assumed in the idealized analytical framework of perfect information.¹ In practice, the use of simple revisable strategies, imitation tactics and rules of thumb seems to be a key ingredient of many decision processes. Thus, when analysing a market and its expected behaviour, it seems valuable to go beyond the perfect-information analysis, and also consider other decision procedures that enjoy greater empirical support and which may be deemed plausible for the context at hand. This point is particularly relevant in markets potentially subject to regulation (e.g. oligopolies) and in situations where the perfect-information theoretical analysis of the social interaction reveals the presence of multiple possible equilibria—as is often the case in indefinitely repeated strategic interactions, including oligopolies in particular. Consequently, several different rules for setting prices or production levels in oligopolies have been analyzed. Bigoni and Fort (2013) provide a recent review of the theoretical and experimental literature on learning in oligopolies.

¹This statement does not necessarily imply that market predictions made using the perfect-information model are irrelevant in real life; the famous “as if” theory of Friedman (Friedman 1953) proves sufficiently accurate and useful in many contexts.

S.S. Izquierdo (✉)

Department of Industrial Organization, Universidad de Valladolid, 47011 Valladolid, Spain
e-mail: segis@eii.uva.es

L.R. Izquierdo

Department of Civil Engineering, Universidad de Burgos, 09001 Burgos, Spain
e-mail: luis@izquierdo.name

In this paper we analyse Cournot oligopolies in which some firms provide a homogeneous good or service and have to choose their production level q_i . We consider that the market process advances in discrete time steps and at every time step the companies have to simultaneously choose whether to increase or decrease the value of their decision variable (q_i). The decision rule considered here can be simply stated as: repeat your last action (i.e. an increase or a decrease in production) if your profits have grown; otherwise, choose the opposite action. This simple rule has been named “Win-Continue, Lose-Reverse” (WCLR) by Huck et al. (2003),² who conducted a thorough study of its convergence properties in symmetric Cournot oligopolies.

The WCLR rule adjusts the level of the decision variable in the direction that is expected to make profits grow, according to the observed effect on profits of the last increment/decrement. Note that this gradual adjustment strategy can be considered a type of reinforcement learning rule: an action (i.e. an increase or decrease in production) is deemed satisfactory—and therefore repeated—if it provides a profit boost, and it is considered unsatisfactory—and therefore avoided—otherwise.

Mathematically, the WCLR strategy presents some similarities with a gradient ascent optimization method. In fact, if the profits of a company were to depend only on its own price or level of production (as in a monopoly with stable demand and costs), this rule would be a gradient ascent method and, under conditions that are well known in the optimization literature (Snyman 2005), it would lead to the vicinity of a local optimum. In a duopoly, however, the profits of a company depend also on its competitor’s price or output level, and the application of the WCLR rule by each of the companies independently does not constitute a gradient ascent method for the joint profit of the two companies. Thus, it is interesting to study to which reference point of the strategic game (e.g. collusive outcome, competitive outcome, or one-shot Nash equilibrium) such a simple strategy converges, if it does converge to any at all.

For a Cournot duopoly in which companies vary their production levels q_i by a predefined amount δ (step size), Huck et al. (2003) show that, under rather general conditions, for small values of δ , the quantities q_i converge to a small area around the cooperative (collusive) solution. In this paper, we show that the convergence of the WCLR rule to collusive outcomes is not robust to small independent perturbations in the profit functions of the firms (e.g., small independent variations in the cost functions, or small differences on the price received by each company). The existence of such small independent perturbations tends to push the process towards the Nash equilibrium of the one-shot game.

The structure of the remaining of the paper is very simple: in Sect. 2 we present the results for the Cournot model, and then we end with the conclusions.

²The same authors use the name “trial and error” in (Huck et al. 2004), where they also present and discuss this learning rule in a discrete-time setup, though the analysis in that paper is focused on a continuous version of the process.

2 Competition in Quantities: Cournot Model

In this section we analyse a Cournot duopoly in which at every time step t ($t = 0, 1, \dots$) each company i ($i = 1, 2$) chooses a production level or quantity $[q_i]_t$. The market price $[p]_t$ is the same for both companies and it depends on the total quantity produced by both firms. The amount $[q_i]_t$ is produced on period t with a cost function $C(q)$. The profit for each company on period t is $[\pi_i]_t = [p]_t [q_i]_t - C([q_i]_t)$. Incremental values are naturally defined as $[\Delta\pi_i]_t := [\pi_i]_t - [\pi_i]_{t-1}$, for $t > 0$, and initial values at time step 0 are $[\Delta\pi_i]_0 = 0$, and $[\Delta q_i]_0 = 0$.

Let us also define $[s_i]_t := \text{sign}([\Delta q_i]_t [\Delta\pi_i]_t)$. Note that $[s_i]_t$ is equal to $+1$ if the last changes in $[q_i]_t$ and $[\pi_i]_t$ took place in the same direction, and $[s_i]_t$ is equal to -1 if such changes went in opposite directions.

For each company i , the production levels are calculated as $[q_i]_{t+1} = \max([q_i]_t + [\Delta q_i]_{t+1}, 0)$, starting with some initial positive production level $[q_i]_0$ at time step 0. The decision rule WCLR used to calculate the production increments $[\Delta q_i]_{t+1}$ is implemented as follows:

WCLR Rule:

- If $t = 0$ or $[s_i]_t = 0$, then $[\Delta q_i]_{t+1}$ takes one random value out of the set $\{-\delta_i, 0, \delta_i\}$, where δ_i is the step size.
- Otherwise, $[\Delta q_i]_{t+1} = \delta_i [s_i]_t$.

It is also assumed that the process includes some “noise” such that, with a small probability ε for each company in every period, the company will deviate from the value prescribed above for $[\Delta q_i]_{t+1}$ and will take a random choice out of the set $\{-\delta_i, 0, \delta_i\}$. This “decision noise” can represent occasional mistakes or experimentation.

Huck et al. (2003) prove that, with $\delta_i = \delta$, under rather general conditions, if the step size δ and the noise level ε are sufficiently small (but strictly positive), in the long run the process $[q_1, q_2]_t$ will spend most of the time in a small neighbourhood around the collusive outcome, and their simulations show a quick convergence to that situation. The remaining of this section is devoted to show that this convergence can be very sensitive to small independent perturbations in the profit functions of the firms. The reader can run all the simulations reported here using the online model at <http://luis.izqui.org/models/wc-lr-cournot/>. The computer model has been implemented in NetLogo (Wilensky 1999).

2.1 The WCLR Rule in the Cournot Model with Noise

For illustrative purposes we consider a linear inverse demand function: $p = \max(100 - (q_1 + q_2), 0)$ and a quadratic cost function: $C(q) = 10q + 0.1q^2$. In this situation, the collusive value for the production of each company, characterized by the first-order conditions $\frac{\partial(\pi_1 + \pi_2)}{\partial q_i} = 0$, is $q_i = 21.43$, which corresponds to

a price level $p = 57.14$. The Cournot equilibrium, characterized by the equations $\frac{\partial \pi_i}{\partial q_i} = 0$, is $q_i = 28.13$, corresponding to a price level $p = 43.75$.

We also set $\delta_i = 0.1$ and $\varepsilon = 0.01$. Initial levels of production $[q_i]_0$ are set randomly in the range $[0, 50]$, but note that the model is ergodic (since $\varepsilon > 0$); thus, its long-run behaviour does not depend on initial conditions.

Departing from the baseline scenario above, we study the sensitivity of the model to three types of noise:

1. “Decision noise”, characterised by the parameter ε , as described above.
2. “Cost noise”, characterised by the parameter ε_{cost} , and implemented by altering each firm’s base cost according to the following formula:

$$c_i = (10q_i + 0.1q_i^2) (1 + \varepsilon_{cost}U_i [-1, 1])$$

where $U_i[-1,1]$ denotes a continuous uniform random variable with range $[-1,1]$.

3. “Price noise”, characterised by the parameter ε_{price} , and implemented by giving each firm i a price p_i according to the following formula:

$$p_i = p (1 + \varepsilon_{price}U_i [-1, 1])$$

where p is the price that corresponds to the total level of output using the inverse demand function. This modified model represents small differences in the price that each company gets for its products, which can be due to a number of different reasons, such as random deviations in the quality of the products of a company with respect to the average quality, different times of arrival at the market (which would allow for some variability in demand), different intermediaries with variable commissions, existence of local markets (which would allow for some variability in price), etc.

Figure 1 below shows a representative run for each of the three types of noise.³ It is clear that in the absence of cost noise or price noise, the WCLR rule leads to the collusive outcome, as shown by Huck et al. (2003). In stark contrast, small independent perturbations in the cost function or in the price function seem to destabilise the collusive outcome and push the simulation towards the Cournot equilibrium. The sensitivity of the model to perturbations in price seems to be greater than the sensitivity to perturbations in cost.

To study this effect rigorously, we conducted a computational experiment where we explored different values of ε , ε_{cost} , and ε_{price} . For each value of these variables we conducted 100 simulation runs, and for each of the runs we computed the average price in the simulation (taken over 10^5 time steps, and neglecting the first 10^4 time steps). Figures 2, 3 and 4 below show the results obtained.

³Note that the simulation runs with “cost noise” or “price noise” have $\varepsilon = 0.01$ too, as prescribed in the baseline scenario.




Fig. 1 Density Histograms of the quantities produced by each firm [q_1, q_2] in one representative simulation run of 100,000 time steps. The *left-most histogram* shows a baseline scenario. The histogram in the *centre* corresponds to a simulation run with a 1 % cost noise added to the baseline scenario, whilst the *right-most histogram* shows a simulation run with a 1 % price noise added to the baseline scenario




Fig. 2 The *blue diamonds* show, for each value of the probability of a random decision ϵ , the mean of 100 prices, each of them obtained from one independent simulation run otherwise parameterised as in the baseline case. The price obtained from each simulation run is the average price in that simulation (taken over 10^5 time steps, and neglecting the first 10^4 time steps). The difference between the minimum average price and the maximum average price across simulations was less than 0.1 in all cases

Figure 2 shows that the WCLR rule leads to collusive outcomes even if the probability of a random decision is fairly high. Figure 3, in contrast, shows that small perturbations in the cost functions of the firms destabilise the collusive outcome and push the process towards the Cournot–Nash equilibrium of the one-shot game. In the same spirit, Fig. 4 shows that the sensitivity of the model to small perturbations in prices is even higher, and the collusive outcome is completely destabilised in favour of the Cournot–Nash equilibrium for values of the price noise as low as 1 %.

Why is the WCLR rule so robust to “decision noise”, but so sensitive to “cost noise” and “price noise”? Note that the stability of the collusive outcome induced by the WCLR rule relies on coordinated moves. When WCLR firms move in the same direction (either increasing or decreasing production levels), they receive




Fig. 3 The *blue diamonds* show, for each value of the cost noise parameter ϵ_{cost} , the mean of 100 prices obtained from 100 independent simulation runs otherwise parameterised as in the baseline. The price obtained from each simulation run is the average price in that simulation (taken over 10^5 time steps, and neglecting the first 10^4 time steps). The *dashed lines* join the minimum average prices and the maximum average prices across simulations




Fig. 4 The *blue diamonds* show, for each value of the price noise parameter ϵ_{price} , the mean of 100 prices obtained from 100 independent simulation runs otherwise parameterised as in the baseline. The price obtained from each simulation run is the average price in that simulation (taken over 10^5 time steps, and neglecting the first 10^4 time steps). The *dashed lines* join the minimum average prices and the maximum average prices across simulations

signals that make them move towards the collusive outcome and linger around it. Alternatively, an uncoordinated move in the vicinity of the collusive equilibrium (possibly due to a perturbation) will make both firms move towards the Cournot–Nash equilibrium in the following time step—assuming no more deviations from the WCLR rule occur. Note, however, that this move towards Nash is itself coordinated, so at the following time step, both firms will simultaneously decrease production and they will keep doing so until they return to the neighbourhood of the collusive outcome. This explains why the collusive outcome is so robust to “decision noise”.

Decision mistakes have an impact only on the decision at the time step at which they occur, and the process goes back towards collusion automatically in two time-steps.

By contrast, the effects of “cost noise” and “price noise” are more profound, as they do not only affect the decision at the time step they occur, but they also have a direct impact on subsequent decisions. This is because these perturbations effectively change the profit landscape and, in that way, they alter the relation between $[\Delta q_i]_t$ and $[\Delta \pi_i]_t$. This deeper type of perturbation, which transcends the time step at which it occurs, is a greater source of miscoordination and, as explained above, uncoordinated moves push the process towards the Cournot–Nash equilibrium.

One final question remains to be answered: why does price variability have a greater impact than cost variability? The answer relates to the different strength with which these two sources of miscoordination affect the profit landscape. It turns out that, given the parameter values used in the illustrations above, profits for both firms in the region of interest are always positive and quite sizable, i.e. income is significantly greater than cost for both firms. In such a situation, a certain percentage change x in prices (and, therefore, in income) induces a greater change in profit than the same percentage change x in costs. Greater changes in profit mean higher chances of altering the sign of $[\Delta \pi_i]_t := [\pi_i]_t - [\pi_i]_{t-1}$, and hence, greater impact on the dynamics of the model. Thus, under such favourable circumstances, it is natural that price variability constitutes a greater source of miscoordination than cost variability. If income and cost were closer in magnitude, the sensitivity of the model to these two types of noise—“price noise” and “cost noise”—would also be more alike. This point can be checked adding a fixed cost equal to 900, i.e. the new cost function reads $C(q) = 900 + 10q + 0.1q^2$. This change makes income and cost similar in the region of interest. In these conditions, the observed impact of cost variability was similar to that of price variability.

2.2 Other Noise Distributions

In this section we show that our results are robust to changes in the noise distribution considered for the price or the cost perturbations. To illustrate this fact, we focus here on a normal distribution with the same mean and standard deviation as the uniform distribution $U[-1,1]$, i.e. the normal distribution $N[0, 1/3]$ with mean 0 and variance 1/3.

First, we show in Fig. 5 below a representative run for each of the three types of noise.⁴ Figure 5, which uses the noise distribution $N[0, 1/3]$ for the price and the cost perturbations, is analogous to Fig. 1, which used the noise distribution $U[-1,1]$.

⁴Note that the simulation runs with “cost noise” or “price noise” have $\varepsilon = 0.01$ too, as prescribed in the baseline scenario.




Fig. 5 Density Histograms of the quantities produced by each firm [q_1, q_2] in one representative simulation run of 100,000 time steps. The *left-most histogram* shows a baseline scenario. The histogram in the *centre* corresponds to a simulation run with a 1 % cost noise following a $N[0, 1/3]$ added to the baseline scenario, whilst the *right-most histogram* shows a simulation run with a 1 % price noise following a $N[0, 1/3]$ added to the baseline scenario




Fig. 6 The *blue diamonds* show, for each value of the cost noise parameter ϵ_{cost} , the mean of 100 prices obtained from 100 independent simulation runs with noise distribution $N[0, 1/3]$, and otherwise parameterised as in the baseline. The price obtained from each simulation run is the average price in that simulation (taken over 10^5 time steps, and neglecting the first 10^4 time steps). The *dashed lines* join the minimum average prices and the maximum average prices across simulations

To study the robustness to changes in the noise distribution, we conducted a computational experiment where we explored different values of ϵ_{cost} and ϵ_{price} using the noise distribution $N[0, 1/3]$, in the same spirit as the experiments shown in Figs. 3 and 4 for noise distribution $U[-1,1]$. Figure 6 below presents the results obtained for ϵ_{cost} , which are very similar to those obtained in Fig. 3. The same similarity was obtained for price perturbations (ϵ_{price}), showing that the sensitivity of the model to cost and price noise does not depend on whether the noise distribution is a uniform distribution or a normal distribution.

2.3 Correlated Perturbations

In this section, we show that the destabilizing factor of the variability in cost or price is not so much the existence of the perturbations, but the fact that they are somewhat independent or uncorrelated between the firms. To illustrate this, here we consider the effect of *correlated* perturbations. Correlations would be observed in the real world if there were variations in costs or in the demand function that affected both companies in a similar way (for instance, seasonal demand variability). To study such situations, we model a price perturbation for each firm which is composed of both a common factor $\alpha U[-1,1]$ —with weight α —and an independent factor $(1 - \alpha) U_i[-1,1]$ —with weight $(1 - \alpha)$ —, according to the formula:

$$p_i = p \left(1 + \varepsilon_{price} R^{\alpha}_i \right)$$

where

$$R^{\alpha}_i = \alpha U[-1, 1] + (1 - \alpha) U_i[-1, 1].$$

Thus, parameter α is a measure of the correlation between the perturbations of each firm. Extreme value $\alpha = 0$ represents completely uncorrelated perturbations (as analyzed above), and extreme value $\alpha = 1$ represents full correlation (where the perturbations for each firm are exactly the same). Figure 7 below shows that the more correlated perturbations are, the less impact they have on destabilising the collusive outcome. As explained before, perturbations affect the dynamics of the model mainly through the generation of miscoordination between the firms; thus, it is natural that the impact of correlated noise, which does not cause so much miscoordination, is less acute than the effect of uncorrelated perturbations.

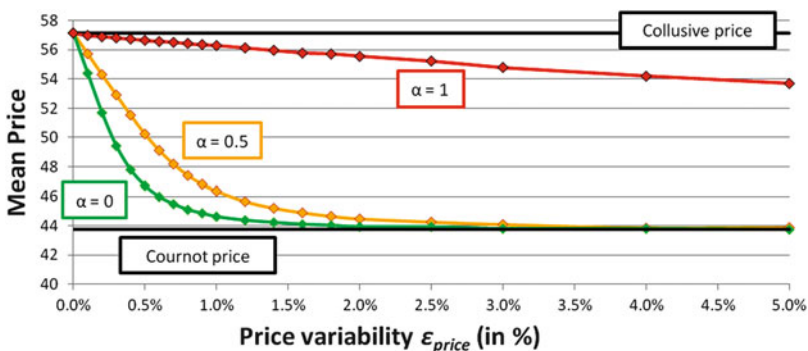


Fig. 7 The diamonds show, for each value of the price noise parameter ε_{price} and different values of α , the mean of 100 prices obtained from 100 independent simulation runs otherwise parameterised as in the baseline. The price obtained from each simulation run is the average price in that simulation (taken over 10^5 time steps, and neglecting the first 10^4 time steps)

2.4 More than Two Competing Firms

The simulation results of Huck et al. (2003) in symmetric oligopolies with more than two competing firms (up to ten) and some small decision noise also showed convergence of the WCLR rule to collusive outcomes. We show in Fig. 8 below that, as in the duopoly case, the existence of small independent perturbations in the price that each company obtains also destabilises the collusive outcome and pushes the process towards the Nash equilibrium of the one-shot game.

Uncorrelated perturbations in cost have the same qualitative effect, so they are not shown here.

It should also be noted that, as the number of competing firms increase, the one-shot Cournot–Nash equilibrium gets closer to the outcome predicted under the assumption of perfect competition, so, as the number of firms increase, the WCLR rule with independent cost or price perturbations leads to market prices and production levels which approach those predicted by the perfect competition theory. Figure 9 below shows the effect of uncorrelated 2 % price perturbations in oligopolies with different number of firms. The results also show an increasing difference between the simulated price and the Cournot price as the number of firms in the market increases, which can be due to the decreasing marginal importance of one firm in the market as the number of firms in the market increases.




Fig. 8 The blue diamonds show, for each value of the price noise parameter ϵ_{price} , the mean of 100 prices obtained from 100 independent simulation runs otherwise parameterised as in the baseline, in an oligopoly with five competing firms. The price obtained from each simulation run is the average price in that simulation (taken over 10^5 time steps, and neglecting the first 10^4 time steps). The dashed lines join the minimum average prices and the maximum average prices across simulations




Fig. 9 The *diamonds* show, for a price noise parameter $\epsilon_{price} = 2\%$ and different number of firms, the mean of 100 prices obtained from 100 independent simulation runs otherwise parameterised as in the baseline. The price obtained from each simulation run is the average price in that simulation (taken over 10^5 time steps, and neglecting the first 10^4 time steps)

Conclusions

The results obtained by Huck et al. (2003) indicate that the simple, individual, “sensible” and not forward-looking decision rule WCLR (“Win-Continue, Lose-Reverse”) can lead to collusion-like outcomes in Cournot oligopolies, even though each company is independently trying to maximize its own profit, and is acting based only on its own past information. Similar results were obtained by Waltman and Kaymak (2008) considering a more involved learning algorithm (Q-learning). In principle, these results could raise important concerns about the fairness of fining firms in oligopolies for apparently carrying out collusive practices, since one could always allege that observed collusion-like outcomes could just be the unintended result of using this type of independent (and thus legitimate) decision rule.

However, this paper has shown that small independent variations in the cost functions, or small uncorrelated perturbations in the price obtained by each firm, can all destabilize the convergence of the WCLR rule to collusive outcomes, pushing the outcomes towards the Nash solution of the one-shot game. Previous simulation results (Keen and Standish 2006) had also indicated that introducing variability in the step sizes used by each company in each period could also push the process towards the Cournot–Nash solution in markets where firms compete in quantities. Consequently, in markets where there is some independent variability over time in the profit functions of the competing firms (which can be due, for instance, to spatially local effects), our results throw doubts on the validity of arguments that try to justify collusion-like outcomes as the unintended result of this kind of “innocent” decision rules.

Acknowledgments We thank three anonymous reviewers for their useful comments. This work has received financial support from the Spanish Ministry of Science and Innovation (CSD2010-00034, SIMULPAST).

References

- Bigoni M, Fort M (2013) Information and learning in oligopoly: an experiment. *Games Econ Behav* 81(1):192–214
- Friedman M (1953) *Essays in positive economics. Part I – the methodology of positive economics.* University of Chicago Press, Chicago, pp 3–43
- Huck S, Normann HT, Oechssler J (2003) Zero-knowledge cooperation in dilemma games. *J Theor Biol* 220(1):47–54
- Huck S, Normann HT, Oechssler J (2004) Through trial and error to collusion. *Int Econ Rev* 45(1):205–224
- Keen S, Standish R (2006) Profit maximization, industry structure, and competition: a critique of neoclassical theory. *Physica A* 370(1):81–85
- Snyman JA (2005) *Practical mathematical optimization: an introduction to basic optimization theory and classical and new gradient-based algorithms.* Springer, Heidelberg
- Waltman L, Kaymak U (2008) Q-learning agents in a Cournot oligopoly model. *J Econ Dyn Control* 32(10):3275–3293
- Wilensky U (1999) *NetLogo.* Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston. <http://ccl.northwestern.edu/netlogo/>