

The Evolution of Walrasian Behavior *

Fernando Vega-Redondo[†]

*I acknowledge financial assistance of the Spanish Ministry of Education, CICYT project no. PB94-1504. I also want to thank the hospitality of the Institute for Economic Development at Boston University and the Economics Dept. of the University of California at San Diego, where this work was completed. Part of the argument used to prove Lemma 2 below was developed jointly with Luis Corchón. I am grateful to him for letting me use it here. I received as well very helpful advice by Jordi Brandts on relevant experimental literature.

[†]University of Alicante and Instituto Valenciano de Investigaciones Económicas.

The Evolution of Walrasian Behavior

Fernando Vega-Redondo

Abstract

Consider an evolutionary context where a given number of quantity-setting oligopolists tend to mimic successful behavior, occasionally experimenting with some small probability. In this context, it is shown that the unique long-run outcome of the process has all firms playing Walrasian, i.e., choosing an output that maximizes profits when taking the market-clearing price as given.

Keywords: Walrasian, Evolution, Cournot, Mutation, Imitation.

1 Introduction

Walrasian Theory builds upon the central hypothesis that agents take prices parametrically, i.e., do not consider the possibility of affecting prices through their consumption or production decisions. In the last decades, there has been a large body of literature whose aim has been to provide some theoretical basis for such Walrasian hypothesis. It can be usefully organized along two separate strands, corresponding to each of the two main areas which are usually taken to integrate classical Game Theory.

On the one hand, there is the approach grounded on Cooperative Game Theory which, motivated by the original work of Edgeworth, revolves around the various so-called Equivalence Theorems. These theorems establish conditions for some of the best known concepts of Cooperative Game Theory (Core, Shapley Value, or the Bargaining Set) to coincide with – or approximate – the set of Walrasian equilibria. (See, for example, Aumann (1964) for an early analysis of Core equivalence, or Hildenbrand (1974) for a systematic review of this issue.)

The second approach to this issue has adopted the perspective of Non-cooperative Game Theory. Here, the focus has been on the exploration of conditions under which *strategic* behavior on the part of producers (and possibly also consumers) will lead to Walrasian behavior. (The seminal work is Gabszewicz & Vial (1972), whereas a good general discussion of this approach can be found in Mas-Colell (1980).)

Quite naturally, in both game-theoretic frameworks – cooperative and noncooperative – a key (necessary) condition for the support of Walrasian behavior is the absence of “monopoly power”, i.e., agents must not have the power to affect prices. In turn, this idea is seen to be intimately linked to the intuitive requirement that the number of agents involved be so large that no particular individual represents more than an insignificant fraction of the whole population.¹

In this paper, an alternative basis is provided for Walrasian behavior which avoids altogether any considerations related to (the absence of) monopoly power, or the related notion of a large enough population. The approach is of an evolutionary nature and, unlike the literature summarized above, explicitly dynamic. In a certain (somewhat distant) sense, it is related to ideas

¹The work of Gabszewicz & Mertens (1971) or Shitovitz (1973) shows that this is not strictly true since the equivalence result can still hold if there are some “atoms” with certain characteristics living in a large world of infinitesimal agents.

discussed long ago by Alchian (1950). In contrast with Friedman (1953) – who strongly argued that selection market forces will always prove effective in selecting for rational behavior – Alchian supported a much more nuanced conclusion, very much in line with our emphasis here. Specifically, he put into question the idea that absolute optimization (or rationality) should always result from the operation of evolutionary forces, stressing instead that it is relative rather than absolute performance which should in the end prove decisive in the long-run.

Two recent pieces of work also related to the present paper are Schaffer (1989) and Rhode and Stegeman (1995). Schaffer has made the following illustrative point.² In a simple context with just *two* quantity-setting firms which have identical and *constant* marginal costs, only Walrasian behavior is evolutionarily stable, i.e., defines an ESS for the corresponding *two-firm* population. (Walrasian behavior, of course, is identified with output choices by both firms for which profits are maximized at the market-clearing price.) On the other hand, Rhode and Stegeman have also arrived to a similar conclusion by applying an stochastic evolutionary approach to a certain two-firm context. Specifically, they show that when two quantity-setting duopolists with quadratic costs face a linear demand curve and imitate the currently most successful output (as well as occasionally “mutate”) the long-run *average* price (and quantities) converge to their Walrasian values.³

This paper shows that the essential gist of these two specialized conclusions does not just apply to their very particular scenarios but is the reflection of a more general state of affairs which can be carried well beyond specific two-firm contexts and a particular cost structure. Specifically, approaching the issue within the dynamic stochastic framework proposed by recent evolutionary literature, Walrasian behavior is shown to evolve in the long run within any quantity-setting oligopoly producing an homogeneous good, provided that the so-called Law of Demand is satisfied (i.e., the market demand curve is downward sloping).

This result seems quite inline with reported evidence on experimental contexts. An interesting case in point is provided by a recent paper by Davis (1995) – see also Holt (1995). Davis studies two alternative scenarios of

²See also the related work by this same author, Schaffer (1988). Another interesting paper which, in the context of coordination games, discusses finite-population effects on evolutionary stability is Crawford (1991).

³Rhode and Stegeman (1995) develop a general approach to the analysis of two-agent contexts, which they then apply to a variety of duopoly contexts (e.g. Bertrand duopolists with differentiated products).

repeated triopoly interaction. One of them is of the Bertrand type, with the (unique) equilibrium of the stage game inducing the Walrasian outcome. A second one is Cournotian, its stage-game equilibrium (also unique) leading to the standard Cournot-Nash outcome (thus, in particular, prices and profits are above the competitive ones).

The experimental conclusions obtained may be summarized as follows. In the first scenario, play settles on the competitive outcome (i.e., the corresponding stage-game equilibrium) in a rather stable manner. In the second one, however, matters are much more unstable, with the average values observed along realized paths of play being significantly closer to the competitive outcome than to the Cournot-Nash equilibrium of the stage game.⁴ Much in line with our emphasis here, the author tentatively interprets these contrasting observations as resulting from players being geared by relative-performance considerations. Whereas these considerations do not interfere with the equilibrium in the first scenario, they significantly do so in the second one (see below for a detailed explanation of these matters).

The model and analysis of the paper can be outlined as follows. Consider a given Cournotian context where an arbitrary (finite) number of firms choose simultaneously their outputs every period. As time proceeds, the most successful behavior (i.e., output choice) tends to spread throughout the market (say, by imitation). Occasionally, with some small probability $\epsilon > 0$, firms also “experiment” (or are simply renewed through some stochastic process of firm turnover). In the limit, as $\epsilon \rightarrow 0$, it is shown that the long-run distribution of the induced stochastic process becomes concentrated in the unique symmetric Walrasian equilibrium. This provides a clear-cut formalization of the idea that, in oligopoly contexts, evolutionary forces will lead firms to behave in a Walrasian fashion, thus behaving “as if” they confronted parametrically the market clearing market price.

A common theme in much of recent Evolutionary Theory revolves around the idea that evolutionary arguments represent a useful way of selecting among different Nash equilibria. (See, for example, Foster and Young (1990), Kandori, Mailath & Rob (1993), or Young (1993).) This has left the impression in some quarters that an evolutionary approach to economic analysis is to be merely viewed as another *equilibrium* selection device, one to be put just along other alternative approaches which can be found in the game-theoretic literature. Even though the original focus on coordination games

⁴Note, in particular, that players do not take advantage of the collusive potential of repeated-game strategies.

has naturally led recent evolutionary literature towards the analysis of equilibrium selection issues, this paper also illustrates the point that evolutionary models may go well beyond these issues. In particular, they may produce interesting behavior which is not a Nash equilibrium.

The rest of the paper is organized as follows. Next section presents the model, Section 3 undertakes the analysis, Section 4 discusses it.

2 The model

Consider a set of firms $N = \{1, 2, \dots, n\}$ involved in a market for an homogeneous product whose demand is summarized by the inverse-demand function $P : \mathbb{R}_+ \rightarrow \mathbb{R}_+$. For every total output $Q \in \mathbb{R}_+$ supplied to the market, this function, assumed decreasing, specifies the market-clearing price $P(Q)$ at which it is sold. All firms are taken to be *ex-ante* symmetric with an identical cost function $C : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ which, for every output q_i produced by any firm $i = 1, 2, \dots, n$, determines its cost of production $C(q_i)$.

For technical reasons, it is assumed that firms have to choose their output from a common finite grid $\Gamma = \{0, \delta, 2\delta, \dots, v\delta\}$, where both $\delta > 0$ and $v \in \mathbb{N}$ are arbitrary. The only condition required is that the Walrasian output \hat{q} belongs to this grid. This output, assumed to exist,⁵ is defined as that which every firm will produce at a market clearing situation when it takes the prevailing price as given, i.e., independent of its output choice. Formally, it is given by the following condition:

$$P(n\hat{q})\hat{q} - C(\hat{q}) \geq P(n\hat{q})q - C(q), \forall q \geq 0. \quad (1)$$

The evolutionary dynamics is taken to proceed in discrete time, which is indexed by $t = 0, 1, 2, \dots$. At each t , the state of the system may be identified with the current output profile $\omega(t) = (q_1(t), q_2(t), \dots, q_n(t))$. Thus, the state space of the system Ω is chosen equal to Γ^n , where Γ is the output grid introduced above. Associated to any such $\omega(t)$, the induced profit profile $\pi(t) = (\pi_1(t), \pi_2(t), \dots, \pi_n(t))$ prevailing at t is defined as follows:

$$\pi_i(t) \equiv P\left(\sum_{j=1}^n q_j(t)\right) q_i(t) - C(q_i(t)), \quad i = 1, 2, \dots, n. \quad (2)$$

At every time t , each firm $i \in N$ is assumed to enjoy a common and independent probability $p > 0$ of being able to revise its former output $q_i(t-1)$. In this event, it is postulated to choose its new output $q_i(t)$ among those which achieved the highest profit in the previous period. More precisely, it is assumed chosen from the set

⁵Standard assumption on costs (for example, non-decreasing marginal costs and small fixed costs) guarantee that a symmetric Walrasian output exists. Provided it exists, the argument used in the proof of the Theorem below ensures that it is unique.

$$B(t-1) = \{q \in \Gamma : \exists j \in N \text{ s.t. } q = q_j(t-1) \& \forall k \in N, \pi_j(t-1) \geq \pi_k(t-1)\},$$

according to a firm-independent probability distribution with full support.⁶

This formulation embodies the customary monotonicity considerations contemplated by evolutionary theory: adjustment dynamics must be responsive to current differential payoffs. As a model of dynamic learning, it may be viewed as a stylized reflection of bounded rationality. Firms, in the real world, live in a very complex environment, in which “imitation of success” could well be a reasonable rule of thumb.⁷ In this respect, it is important to emphasize that our very simple model of the market environment does not intend to mimic the real world but is just a tractable “metaphor” of it. Therefore, firms’ behavioral rules should *not* be discarded merely on the grounds that they are too simplistic (or suboptimal) relative to some “obvious” features of the postulated theoretical framework.

Once firms have adjusted their output as described, firms are also assumed to “mutate,” at every t , with some common independent probability $\epsilon > 0$. In this event, they choose some arbitrary output in Γ , all of them selected with some given positive probability. The interpretation here is that with small probability firms either experiment with new choices or they are replaced by some newcomer which chooses its output from *tabula rasa*.

The Markov process described is clearly ergodic. Specifically, it is positively irreducible since, through the mechanism of mutation, every two states in Ω are directly connected with positive probability. By a standard result in the Theory of Markov Chains, it may be ensured to have a unique invariant

⁶This formulation is much narrower than necessary, and is adopted here for the sake of expositional simplicity. When $B(t-1)$ is not a singleton, it seems reasonable to allow for the possibility that any revising firm whose former output $q_i(t-1)$ lies in this set might continue choosing it with very high probability. If this probability remains short of one (although arbitrarily high), it is wholly consistent with the full-support condition contemplated above. But even in the extreme case where such inertia is assumed complete (i.e., the probability of remaining with the previous output is postulated to be one under those circumstances), the statement of the Theorem can be shown to remain fully valid without modification. The argument, however, becomes somewhat more complicated, as outlined in Footnote 9 below.

⁷As suggested by an anonymous referee, this rule seems intuitively appealing if the decision makers are rewarded according to their relative performance. (For example, managers may be promoted depending on whether they earn a profit at least as large as that of competitors.) Along these lines, note that expression (8) below can be used to establish that, if firms’ payoff functions are identified with relative profits (say, with their own profit minus the average), the symmetric Walrasian equilibrium defines a *Nash equilibrium* of the induced game.

distribution which fully summarizes the long-run performance of the process (in particular, it determines the long-run frequencies with which every state is observed, with probability one, along any sample path). Since this invariant distribution obviously depends on ϵ (the mutation probability), it is denoted by $\mu_\epsilon \in \Delta(\Omega)$, where $\Delta(\Omega)$ is the set of probability measures on Ω .

3 Analysis

Intuitively, one wants to conceive of the mutation probability ϵ as being small. To capture this idea, the analysis will focus on the behavior of the process as ϵ becomes small; or more formally, on the limit invariant distribution $\mu^* \equiv \lim_{\epsilon \rightarrow 0} \mu_\epsilon$. (This limit distribution is easily seen to be a well defined element of $\Delta(\Omega)$ – see below.)

Any state in Ω which belongs to the support of μ^* is usually called a *stochastically stable state*. Due to the ergodicity of the process, only those states which are stochastically stable will be observed a significant fraction of the time, a.s., along any sample path of the process if the mutation probability is small. The following result singles out the unique stochastically stable state of the process.

Theorem *Let $\omega_{\hat{q}} = (\hat{q}, \hat{q}, \dots, \hat{q})$. Then, $\mu^*(\omega_{\hat{q}}) = 1$.*

Proof. As customary in recent evolutionary literature, the proof shall rely on the graph-theoretic techniques developed by Freidlin & Wentzel (1984) and first applied to in the evolutionary literature by Foster & Young (1990). Particularized to our context, they may be briefly summarized as follows.

For each $\omega \in \Omega$, define a ω -tree H as a collection of ordered pairs – “arrows” – (ω', ω'') such that:

- (i) every $\omega' \in \Omega \setminus \{\omega\}$ is the first element of exactly one pair, and
- (ii) from every $\omega' \in \Omega \setminus \{\omega\}$ there exists a *path* $\{(\omega^0, \omega^1), (\omega^1, \omega^2), \dots, (\omega^{s-1}, \omega^s)\}$ such that $\omega^0 = \omega'$ and $\omega^s = \omega$. The set of all such ω -trees is denoted by \mathcal{H}_ω .

Let T_ϵ stand for the transition matrix of the postulated evolutionary dynamics when the mutation probability is ϵ . Define, for each $\omega \in \Omega$,

$$r(\omega) \equiv \sum_{H \in \mathcal{H}_\omega} \prod_{(\omega', \omega'') \in H} T_\epsilon(\omega', \omega''). \quad (3)$$

Then, as established by Freidlin & Wentzel (1984), we have:

$$\mu_\epsilon(\omega) = \frac{r(\omega)}{\sum_{\omega' \in \Omega} r(\omega')} . \quad (4)$$

Each $r(\omega)$ is a polynomial in ϵ . Thus, the *limit invariant distribution* defined above is well-defined and, therefore, unique. To compute each $r(\omega)$, it is useful to introduce a cost function on possible transitions

$$c : \Omega \times \Omega \rightarrow \mathbb{N} \cup \{0\} ,$$

which for each pair (ω, ω') specifies the minimum number of mutations $c(\omega, \omega')$ needed for the transition to occur with ensuing positive probability *via* mutation-free dynamics. That is, if $d(\omega, \omega')$ denotes the number of firms whose output differs between ω and ω' and T_0 stands for the transition matrix corresponding to the mutation-free dynamics, then⁸

$$c(\omega, \omega') \equiv \min_{\omega'' \in \Omega} \{d(\omega, \omega'') : T_0(\omega'', \omega') > 0\}.$$

The function $c(\cdot)$ may be extended to every path h and every tree H , by simply adding the cost of all their constituent links. In view of (3), the order of each $r(\omega)$, as a polynomial in ϵ , is simply given by the minimum number of mutations required along some ω -tree, i.e., $\min_{H \in \mathcal{H}_\omega} c(H)$. Thus, from (4), it follows that the set of stochastically stable states are precisely those whose minimum cost trees are themselves minimum across all possible states in Ω .

As customary, a set $A \subset \Omega$ is defined to be a *limit (or absorbing) set* of the mutation-free dynamics if this set is closed under finite chain of iterations of T_0 . That is:

- (i) $\forall \omega \in A, \forall \omega' \notin A, T_0(\omega, \omega') = 0$.
- (ii) $\forall \omega, \omega' \in A, \exists m \in \mathbb{N}$ s.t. $T_0^{(m)}(\omega, \omega') > 0$.

Let \mathcal{A} be the collection of limit sets of the mutation-free dynamics and denote $a \equiv |\mathcal{A}|$. Clearly, only states which belong to some limit set can be stochastically stable. (All other states are just transitory in the mutation-free dynamics and, therefore, very “infrequent” when the mutation probability becomes infinitesimal.)

Given any particular output $q \in \Gamma$, let $\omega_q = (q, q, \dots, q)$ represent the state where *all* firms choose this output. Such states will be called *monomorphic*. Since, obviously, every monomorphic state defines a singleton limit set of T_0 , the proof of the theorem follows directly from the following lemmata.⁹

⁸Note that, for the sake of formal simplicity in the argument, the cost function defined is slightly different from that introduced by, say, Kandori, Mailath & Rob (1993).

⁹Lemmas 1 and 2 would still apply if (violating the full-support requirement postulated in Section 2) the adjustment dynamics were to display full inertia when the firm’s output was one of those formerly leading to highest profits. Under these circumstances, the main complication in the argument derives from the fact that some limit states of the mutation-free dynamics may be polymorphic (i.e., *not* monomorphic). But even in this case, the essential point to notice is that any $\{\omega\} \in \mathcal{A}$ different from the symmetric Walrasian state can still be de-stabilized with just one mutation. If ω is monomorphic, the argument is as in Lemma 1. If it is polymorphic, consider a transition to some other state ω' where one of the firms previously choosing q mutates to q' , a different output chosen at ω by some other firm. Without relying on any further mutation, an ensuing transition to some other stationary $\omega'' \neq \omega$ becomes then possible. If ω' is a limit state of the mutation-free dynamics, this is achieved trivially. Otherwise, it must happen that *not* all outputs

Lemma 1 *There exists an $\omega_{\hat{q}}$ -tree $\hat{H} \in \mathcal{H}_{\omega_{\hat{q}}}$ such that $c(\hat{H}) = a - 1$.*

Lemma 2 *Let $A \neq \{\omega_{\hat{q}}\}$ be a limit set of the mutation-free dynamics. $\forall \tilde{\omega} \in A$ and every $\tilde{\omega}$ -tree $\tilde{H} \in \mathcal{H}_{\tilde{\omega}}$, it follows that $c(\tilde{H}) \geq a$.*

Proof of Lemma 1. First, notice that a set $A \subset \Omega$ is a limit set of the mutation-free dynamics if, and only if, it is a singleton consisting of a monomorphic state. (Thus, \mathcal{A} is isomorphic to the set Γ .) The “if” part is obvious from the specification of the process (i.e., if the state is monomorphic, every strategy revision possibility by any firm will leave the state unchanged). On the other hand, the “only if” follows from the fact that strategy revision is a firm-independent phenomenon whose probability density at each t is assumed to have full support on the respective set $B(t-1)$. Therefore, there is always positive probability (bounded above zero, since the state space is finite) that all firms adjust their strategy towards the same output.

Consider any $q \in \Gamma$ such that $q \neq \hat{q}$. It is next shown that $c(\omega_q, \omega_{\hat{q}}) = 1$. To verify this claim, it is enough to show that when all firms except one produce q , the firm that produces \hat{q} obtains strictly higher profits than the rest. That is,

$$P((n-1)q + \hat{q})\hat{q} - C(\hat{q}) > P((n-1)q + \hat{q})q - C(q) \quad (5)$$

First, it is argued that

$$[P(n\hat{q}) - P((n-1)q + \hat{q})]\hat{q} < [P(n\hat{q}) - P((n-1)q + \hat{q})]q. \quad (6)$$

If $q < \hat{q}$, $[P(n\hat{q}) - P((n-1)q + \hat{q})]$ is negative since $P(\cdot)$ is decreasing, and the above inequality obviously follows. If $q > \hat{q}$ instead, $[P(n\hat{q}) - P((n-1)q + \hat{q})] > 0$ but (6) is still preserved.

Now, rewrite (6) as follows:

$$P(n\hat{q})\hat{q} + P((n-1)q + \hat{q})q < P(n\hat{q})q + P((n-1)q + \hat{q})\hat{q}.$$

Subtracting the term $[C(q) + C(\hat{q})]$ from both sides of the above expression we obtain:

$$\begin{aligned} [P(n\hat{q})\hat{q} - C(\hat{q})] + [P((n-1)q + \hat{q})q - C(q)] < \\ [P(n\hat{q})q - C(q)] + [P((n-1)q + \hat{q})\hat{q} - C(\hat{q})] \end{aligned} \quad (7)$$

chosen at ω' lead to the same profit. Thus, with positive probability, some output choice present at ω will fully disappear from the population. If this happens, the stationary point eventually attained by the process (without mutation) must be different from ω . Combining this idea with the fact that the Walrasian state still needs two mutations to become “irreversibly de-stabilized” (cf. Lemma 2), an analogous line of proof can be used to complete the argument.

From (1), the first term in the LHS of (7) is no smaller than the first term in its RHS. Therefore, one must have that the second term in the LHS of this expression is strictly smaller than its second term in the RHS. But this is just what (5) expresses, which confirms the desired claim.

It is now verified that there is $\omega_{\hat{q}}$ -tree \hat{H} whose cost $c(\hat{H}) = a - 1 = v$, where recall that $v + 1$ is the (arbitrary) cardinality of the set Γ . To construct such a tree consider first the v links $\{(\omega_q, \omega_{\hat{q}}) : q \neq \hat{q}\}$. The aggregate cost of these links is v . But then, since monomorphic states are the only limit sets of the mutation-free dynamics (recall above), the remaining states can be linked to them in a costless manner to complete a full $\omega_{\hat{q}}$ -tree whose total cost is v . This completes the proof of the Lemma.

Proof of Lemma 2. Let $\{\omega_{\tilde{q}}\}$, $\tilde{q} \neq \hat{q}$, be some limit state of the mutation-free dynamics. Every $\tilde{\omega}$ -tree \tilde{H} must incur a cost $c(\tilde{H}) \geq v$, since at least one mutation is needed to escape *every* one of the v limit sets $\{\omega_q : q \neq \tilde{q}\}$. But in fact this lower bound can be chosen equal to $v + 1$ since, as presently shown, at least two mutations are necessary to escape the limit set $\{\omega_{\hat{q}}\}$. To show this, it must be confirmed that with just one mutation from state $\omega_{\hat{q}}$, Walrasian firms still earn a profit higher than the non-Walrasian firm. That is, for all $q \neq \hat{q}$,

$$P((n-1)\hat{q} + q)\hat{q} - C(\hat{q}) > P((n-1)\hat{q} + q)q - C(q). \quad (8)$$

Again, the fact that $P(\cdot)$ is decreasing implies that:

$$[P(n\hat{q}) - P((n-1)\hat{q} + q)]\hat{q} < [P(n\hat{q}) - P((n-1)\hat{q} + q)]q.$$

Hence subtracting the term $[C(q) + C(\hat{q})]$ from both sides and re-arranging terms one obtains:

$$\begin{aligned} [P(n\hat{q})\hat{q} - C(\hat{q})] + [P((n-1)\hat{q} + q)q - C(q)] < \\ [P(n\hat{q})q - C(q)] + [P((n-1)\hat{q} + q)\hat{q} - C(\hat{q})] \end{aligned}$$

which, relying on (1), implies (8), as claimed.

Combining Lemmas 1 and 2, the proof of the theorem is complete. ■

Remark 1 *As explained in Lemma 1, only one (suitable) mutation is required to trigger a transition to the Walrasian monomorphic state $\omega_{\hat{q}}$ from any other stationary state. This implies that the rate at which the stochastic process will converge to the long-run outcome will be relatively fast (i.e., of order ϵ , the mutation rate). In particular, it is independent of population size, in contrast with what is often the case in recent evolutionary models. (Ellison (1993) is a well-known exception. However, notice that, unlike in his case, fast convergence here does not depend on interaction being local.)*

4 Discussion

The above result starkly reflects the essential differences between the evolutionary and classical approaches to the analysis of game-theoretic situations. In the classical approach, a (Nash) equilibrium is defined as a strategic configuration where no player can unilaterally deviate and become better off in *absolute* terms, i.e., without any concern for its ensuing relative standing *vis a vis* the other players. In contrast, evolutionary models rely exclusively on relative payoff magnitudes to define equilibrium (or stationary) situations. In any such situation, the evolutionary approach requires that no player obtains a payoff lower than some of its opponents. In other words, in the logic of evolutionary theory, differential payoffs is all that matters in order to evaluate the long-run feasibility of a given configuration. In Biology, this is rooted in the Darwinian theme of individuals with differential reproduction rates struggling to prevail in a finite-capacity environment. In the social contexts, it could also reflect ideas of differential survival *stricto sensu*, but in general it is best viewed as responding to forces of learning and imitation.

This “concern” for relative payoffs is what some biologists (see Hamilton (1970)) have called *evolutionary spite*. Due to these considerations, it may be “evolutionarily sound” to undertake deviations from the status quo which will worsen one’s own payoff if they also decrease even further that of the opponents. This is precisely what underlies Lemma 1 above. When a firm deviates from a monomorphic non-Walrasian configuration to producing the Walrasian output, it may well decrease its payoff in absolute terms (for example, if the original state was the Cournot-Nash equilibrium). However, it will always hurt its competitors even further. Reciprocally, Lemma 2 builds upon the idea that from a monomorphic Walrasian configuration, no single firm can ever deviate and become relatively better off.

Of course, for these considerations to be at all relevant the market must involve a finite number of firms (no matter how large). However, if each given firm was truly insignificant relative to the market, spiteful behavior would be pointless. In the limit, as one approaches a context with a progressively larger number of firms, the room for spiteful considerations also shrinks and, consequently, it is easy to see that the long-run outcome would approximate the Cournot-Nash equilibria of the game. This is just a manifestation of the well-know result that, as the number of firms expands in an oligopolistic context, the Cournot-Nash equilibria converge to the Walrasian outcome. In our case, however, this outcome is achieved for any finite number of firms.

Even though the number of active firms is a given parameter which does not affect our qualitative conclusions, it would be a natural and interesting extension to have the population size become an *endogenous* variable of the

model. In general, firm survival, entry, and exit, must represent important ingredients of any evolutionary approach to the study of industrial dynamics. Such population adjustments should also be instrumental in having firms' characteristics (e.g., their costs) become endogenously determined by the evolutionary process. In a sense, the simple model analyzed here may be viewed as singling out some of the considerations which would be likely to play a significant role in the "quantity dimension" of such a richer framework.

Let us conclude with a brief re-consideration of one the questions which initiated the paper: Is it true that, as Friedman (1953) claimed, evolutionary forces lead to as-if rational agents? In our context, the answer has to be in the negative if rationality is identified with behaving optimally in the short-run, given full knowledge of the underlying circumstances. In a finite oligopoly, Walrasian behavior is not optimal if firms are aware of the characteristics of the game and attempt to maximize their instantaneous payoff. However, if firms' objective function is taken to include survival as a primary consideration (for example, if their intertemporal discount rate is close to one) and their survival is linked to relative payoffs (e.g., the larger is the accumulated profit relative to that of the competitors, the stronger the firm is to launch and/or repel a predatory campaign), it may well be that the rationality of Walrasian behavior could be recovered. A detailed exploration of these issues in the context of bankruptcy games (see, for example, Rosenthal & Rubinstein (1984)) is a task left for future research

References

- [1] Alchian, A.A. (1950): "Uncertainty, evolution, and economic theory," *Journal of Political Economy* **57**, 211-21.
- [2] Aumann, R. (1964): "Markets with a continuum of traders," *Econometrica* **32**, 39-50.
- [3] Crawford, V. (1991): "An 'evolutionary' interpretation of van Huyck, Battalio, and Beil's experimental results on coordination," *Games and Economic Behavior* **3**, 25-59.
- [4] Davis, D.D. (1995): "Advance production, Cournot outcomes and trigger strategies: an experimental investigation," Virginia Commonwealth University.
- [5] Ellison, G. (1993): "Learning, local interaction, and coordination," *Econometrica* **61** (1993), 1047-62.
- [6] Foster, D. & P. Young (1990): "Stochastic evolutionary game dynamics," *Theoretical Population Biology* **38**, 219-32.
- [7] Freidlin, M.I. & A.D. Wentzel (1984). *Random Perturbations of Dynamical Systems*, New York: Springer Verlag.
- [8] Friedman, M. (1953): *Essays in Positive Economics*, University of Chicago Press, Chicago.
- [9] Gabszewicz, J.J. & J.-P. Vial (1972): "Oligopoly à la Cournot in a general equilibrium analysis," *Journal of Economic Theory* **4**, 381-400.
- [10] Gabszewicz, J.J. & F. Mertens (1971): "An equivalence theorem for the core of an economy whose atoms are not too big," *Econometrica* **39**, 713-21.
- [11] Hamilton, W.D. (1970): "Selfish and spiteful behavior in an evolutionary model," *Nature* **228**, 1218-20.
- [12] Hildenbrand, W. (1974): *Core and Equilibria of a Large Economy*, Princeton Univ. Press, Princeton.

- [13] Holt, C.A. (1995): "Industrial Organization: A survey of the results of laboratory experiments," in J Kagel & A. Roth, eds., *The Handbook of Experimental Economics*, Princeton Univ. Press, Princeton.
- [14] Kandori, M., G. Mailath & R. Rob (1993): "Learning, mutations, and long-run equilibria in games," *Econometrica* **61**, 29-56.
- [15] Mas-Colell, A. (1980): "Non-cooperative approaches to the theory of perfect competition," *Journal of Economic Theory* **22**, 121-376.
- [16] Rhode, P. and M. Stegeman (1995): "Non-Nash equilibria of Darwinian dynamics (with applications to duopoly)," mimeo.
- [17] Rosenthal, R. & A. Rubinstein (1984): "Repeated two-player games with ruin," *International Journal of Game Theory* **13**, 155-77.
- [18] Schaffer, M.(1989): "Are profit maximizers the best survivors," *Journal of Economic Behavior and Organization* **12**, 29-45.
- [19] Schaffer, M.(1988): "Evolutionarily Stable Strategies for a Finite Population and a Variable Contest Size," *Journal of Theoretical Biology* **132**, 469-78.
- [20] Shitovitz, B. (1973): "Oligopoly in markets with a continuum of traders," *Econometrica* **41**, 467-502.
- [21] Young, P. (1993): "The evolution of conventions," *Econometrica* **61**, 57-84.