Società italiana di economia pubblica

XIV CONFERENZA

IL FUTURO DEI SISTEMI DI
WELFARE NAZIONALI TRA INTEGRAZIONE EUROPEA E DECENTRAMENTO REGIONALE
coordinamento, competizione, mobilità
Pavia, Università, 4 - 5 ottobre 2002

THE RATIONALITY OF INFORMATION GATHERING: OLIGOPOLY

ISABELLA IMPERATO

pubblicazione internet realizzata con contributo della

COMPAGNIA di San Paolo

società italiana di economia pubblica

dipartimento di economia pubblica e territoriale – università di Pavia

First draft
14 September 2002

THE RATIONALITY OF INFORMATION GATHERING: OLIGOPOLY

By Isabella Imperato

1. Introduction

This paper is concerned with the optimal use of information in a standard oligopoly model, where information processing is costly. The basic motivation for this work comes from the observation that often expectations are not perfectly (or not at all) correlated to fundamentals, and that the "mood of the market" can change abruptly, even in the absence of news. In the existing literature, the most appealing explanations for these phenomena focus on the mechanism of coordination of beliefs and the revelation of insider information through the process of trading. A certain degree of informational asymmetry among economic agents is often exogenously given. However, such assumptions can be applied only to some well-specified markets, notably the financial markets. Other attempts to throw sand in the wheels of the basic rational expectations hypothesis (REH) comprise approaches where sluggish action is the result of either institutional constraints, or costs related to action, and approaches that modify the expectation process itself. The two approaches are generally seen as unrelated, yet this dichotomy may be too extreme.

The REH contains an internal contradiction, first underlined Grossman and Stiglitz (1980): if everyone has free access to information, nobody is rewarded for the real resources that go into its collection and processing, so that an "informationally efficient" equilibrium does not exist. In the same vein, this paper shows how explicit analysis of the cost of processing information can lead to optimally designed time-dependant rules, where there are periods during which a rational agent chooses to be less than fully informed, so that the property of orthogonality among forecast errors and available information does not necessarily hold: learning occurs at discrete intervals.

I apply the above concepts to a simple oligopoly model, where agents must pay a fixed, sunk cost for using information about the state. This cost is referred to as the cost of *thinking,* intending the procedure through which people organise and then elaborate information. The first result of the paper is that *thinking* occurs at discrete time intervals: the optimal lag between each subsequent updating of the information set is increasing in the cost of *thinking* itself (information need not be costly here, it is sufficient that it is costly to process it), while it is decreasing in the unconditional variance of the stochastic state variable (bigger uncertainty produces more intensive *thinking*). The second result – and possibly the more interesting from the point of view of public policy - is that, even if the producers do not share information, an informational externality is present, so that firms will optimally *think* at different times: *thinking* is staggered, and asymmetric information becomes an endogenous feature of the equilibrium – consistently with Grossman and Stiglitz' (1980) results. Finally, I find a negative correlation between the incentive for the firm to acquire information and the degree of competition on the market: monopoly and oligopoly firms use more information than competitive firms.

The interactions between the degree of market competition and the "information intensity" of an economy may have interesting (and disregarded) implications for the design of public policies. Think for example at market liberalisation policies: it may well be the case that induced differences in the frequency of decision-making represent a theoretical case for accompanying liberalizations with policies of disclosure and transparency, aimed at reducing information costs.

2. The model

The market is constituted by n identical firms, each of them maximising her revenues. Demand at time t is represented by the following price equation

$$p(t) = \alpha(t) - \beta q(t) \tag{1}$$

where p denotes price and q is quantity. Overall production q is the sum of individuals productions q_i. The intercept α is stochastic: it obeys a mean reverting process

$$d\alpha = k(\mu - \alpha)dt + \sigma dz \tag{2}$$

with long-run mean μ, speed of adjustment k>0, and diffusion coefficient σ; z is a standard Wiener process. The unconditional mean of the process is μ, and its unconditional variance is equal to $\sigma^2/2k$. Given an initial condition

$$\alpha(\tau) = \alpha_\tau \qquad \tau \le t \tag{3}$$

the conditional first and second moments of the process can be simply derived by multiplying both sides of (2) by exp(kt) and integrating by parts between τ and t, getting

$$E[\alpha(t) \mid \alpha_\tau] = \mu - (\mu - \alpha_\tau)e^{-k(t-\tau)}$$

$$V = \mathrm{var}[\alpha(t)|\alpha_\tau] = E\left(\{\alpha(t) - E[\alpha(t)|\alpha_\tau]\}^2\right) = \frac{\sigma^2}{2k}\left[1 - e^{-2k(t-\tau)}\right] \tag{4}$$

For the sake of simplicity, production costs are normalised to zero. I hence assume that output is chosen by each producer to maximise her revenues, net of any other costs that may occur in the process. It is well known that repeated interaction among oligopolists yields multiple equilibria. The aim of the present paper, however, is not to add any insights into the Cournot repeated game problem, but to introduce in oligopoly a trade-off between "good forecasts" and their cost. For this purpose, I will pick up only one out of the possible equilibria, that is the one-shot Nash solution reproduced at each stage of the game.

It now becomes crucial to specify the informational characteristics of the present model. First, I rule out model uncertainty by assuming that both the structure of the model and the stochastic process, which governs the evolution of the intercept of the demand function, are known. Moreover, I assume that the information about the value taken by contemporaneous α is freely available, but that using it to make an output choice entails a fixed, sunk cost equal to b, to which I shall henceforth refer to as the *cost of thinking*.[1] *Thinking* is the procedure through which people organise and then elaborate information. The agent has to devise a time-dependant strategy where she decides how frequently to *think*, and how frequently to pay the related cost. The latter can hence be interpreted in many ways. It could be the fixed cost of accessing facilities that will help with the decision process, such as libraries, consultants, time. It could also be the effort required to switch from a routinely behaviour to proactive – and stressful – assessment of whether change is required. In any case, it should not be confused with the cost of information itself: it is specific to the agent and her environment, not to the type or amount of information needed.

When the producer *thinks*, she acquires full information on all past and current values of prices, quantities and shocks. On the contrary, she cannot directly observe whether her opponents are as well *thinking*. When *non-thinking*, the oligopolist does not acquire or use new information, so

[1] This concept has been already introduced in a paper co-authored with David K.H. Begg on "The rationality of Information gathering: Monopoly", in The Manchester School, Vol. 69 n. 3, 237-252, June 2001.

that she must rely on information acquired in the past (at $\tau_i \leq t$ for example) and update it in a "mechanical" way, which takes into account the structure of the model but not the shocks that may have taken place lately (from τ_i to t).

At each stage of the game, each producer decides whether or not to *think*, then she fixes production at her best guess of the "Cournot" level. With full information, the quantity supplied at any time by agent i, given the other agents' productions, would be

$$q_{it}^{*} = \frac{\alpha_t - \beta \sum_{j \neq i} q_{jt}}{2\beta} \tag{5}$$

However, contemporaneous α cannot be taken into account unless the producer pays a cost equal to b, i.e. *thinks*. Hence, I must allow for some degree of "ignorance". At the same time, I will allow for some degree of asymmetric information, since there is no guarantee that producers will *think* together. In the following, I define $\tau_1,\ldots,\tau_n$ the latest *thinking* times of producers from 1 to n, where $\tau_1 \leq, \ldots, \leq \tau_n$, so that producer n is the best informed. Her effective reaction function is

$$q_{nt} = \frac{1}{2\beta}\left[E_n(\alpha_t) - \beta \sum_{j=1}^{n-1} q_{jt}\right] \tag{6}$$

where to avoid cumbersome notation I term E_n the expectation conditional on having last thought at time τ_n. Producer n-1 is more informed than the others but producer n: her reaction function is

$$q_{n-1,t} = \frac{1}{2\beta}\left[E_{n-1}(\alpha_t) - \beta \sum_{j=1}^{n-2} q_{jt} - \beta E_{n-1}(q_{nt})\right] \tag{7}$$

By substituting the conditional expected value of (6) into (7) I get

$$q_{n-1,t} = \frac{1}{3\beta}\left[E_{n-1}(\alpha_t) - \beta \sum_{j=1}^{n-2} q_{jt}\right] \tag{8}$$

Applying recursively this method to all agents, I can write down a whole system of reaction function, whose generic one is

$$q_{n-i,t} = \frac{1}{(i+2)\beta}\left[E_{n-i}(\alpha_t) - \beta\sum_{j=1}^{n-i-1} q_{jt}\right] \tag{9}$$

System (9) is a system of n equations, where the quantity produced by a single agent q_i depends on that produced by the others and on the producer's expected α value. After solving, each q_i will depend on the α's expected by each of the producers, that is on all producers' information sets. Given that information sets depend on the τ's – the last *thinking* times of the producers – production levels are as well functions of the τ's, and are specified for any given set of *thinking* times. In other words, I can interpret the system (9) as a collection of reaction functions in the *thinking* times τ. The τ's are uncertain variables, but in a strategic sense.[2]

Equations (1) and (5) together imply that – given the competitors' production levels – individual instantaneous revenues are smaller than full information revenues by an amount proportional to the quadratic difference between full information and effective production $(q_i^*-q_i)^2$, which in the following I will label F. The latter can thus be interpreted as the cost of not taking into account all the available information. Since the *thinking* cost is equal to b, the trade-off between *thinking* and *non-thinking*, or between making an active decision and continuing the preceding routine, will depend on the instantaneous loss:

$$E\left[(q_i^* - q_i)^2\right] + b = F + b \tag{10}$$

where $b \geq 0$ in a *thinking* time and $b=0$ otherwise.[3]

3. The optimisation

The relevant instantaneous loss for the optimisation is that of firm 1.[4] Indeed, given the above assumptions on the informational structure, producer 1 is either as well informed as the others, and her case is representative, or she is the less informed, and hence she is the one who is in the process of deciding whether or not to *think*.[5] In other terms, whenever the *thinking* time comes,

[2] In the same way as in a textbook Cournot model there is strategic uncertainty about one's opponent's production level, which is consequently written in levels and not in expectations inside the reaction function.
[3] For a diagrammatic illustration of the above concept, see Begg and Imperato (2001), section 2. See the same paper for a throughout discussion on the implication of the linearity assumption of the demand curve.
[4] In the following, I will shift to the appendices the algebra that was already developed – or is very similar to that shown – in Begg and Imperato (2001).
[5] I have excluded the possibility of some herd behaviour, where a portion of firms *think* contemporaneously at some dates and the others either contemporaneously at some other dates, or as above at different rates. This because- as will

the firm who is going to *think* is the less informed, and this is exactly the point where a balance must be found between the cost of *thinking* and its revenue (that is, the cost of *non-thinking*).

Applying equations (5) and (9) to the first producer, after some algebra[6] the expected squared deviation of her production from its full information level turns out to be equal to

$$E[(q_{1t}^{*}-q_{1t})^{2}]=\frac{\sigma^{2}}{4\beta^{2}}\frac{1}{2k}\{1-e^{-2k(t-\tau_{n})}\frac{e^{-2k(t-\tau_{n})}}{4}[1-e^{-2k(\tau_{n}-\tau_{n-1})}]+ \\ +\frac{e^{-2k(t-\tau_{n-1})}}{9}[1-e^{-2k(\tau_{n-1}-\tau_{n-2})}]+...+\frac{e^{-2k(t-\tau_{2})}}{n^{2}}[1-e^{-2k(\tau_{2}-\tau_{1})}]\} \tag{11}$$

The oligopolist faces a problem of optimal stopping: she must decide at every time t whether to apply the control (*think*) and get some termination payoff, including the cost b of exercising the control, or to continue to rely on some old information. The variable under control is the expected squared deviation of output from its full information level q* (F). This variable is left free to rise until it reaches the upper border of the continuation region, then it is driven back to zero by *thinking*. Hence, agents must devise a time-dependant optimal strategy defined in terms of *thinking* times. The optimisation implies exploitation of the trade-off between the cost of *thinking* (b) and its benefit (a smaller expected squared deviation between effective and full information production). In the following, I shall define Δ the time lag between the last thinking time τ and the next; at t-τ=Δ the state variable F reaches its upper barrier, that I shall call m.

3.1 Duopoly

To make understanding easier, I will begin with duopoly. With n=2, equation (11) becomes

$$E\left[\left(q_{t1}^{*}-q_{t1}\right)^{2}\right]=\frac{\sigma^{2}}{4\beta^{2}}\frac{1}{2k}\left\{1-e^{-2k(t-\tau_{2})}+\frac{e^{-2k(t-\tau_{2})}}{4}\left[1-e^{-2k(\tau_{2}-\tau_{1})}\right]\right\} \tag{12}$$

The global loss of producer 1 over the infinite time horizon is given by the present discounted value of losses F plus *thinking* costs:

be shown in the following, *thinking* all together is not an equilibrium, and the same reasoning that proves the above statement can be used to rule out the appearance of *thinking groups*. This is of course a consequence of excluding the possibility of collusive strategies.

[6] See Appendix 1.

$$L = \int_{\tau_1}^{\infty} F e^{-\rho(s-\tau_1)} ds + \mathit{thinking}\ \cos ts \tag{13}$$

where $0<\rho<1$ is the discount rate.

Graphically, we have

```
          τ1                    τ2          t            τ1+Δ
---------+--------------------+--------------------+----------
```

The relevant continuation region is the interval (τ_1, $\tau_1+\Delta$). At the upper border of this region, producer 1 *thinks*, that is she pays the cost b and brings q_1*-q_1 to zero.

In the following, I will rely on the methodology shown in Dixit (1991), where a similar problem is treated by means of its discrete approximation. For an intuitive treatment, see also Dixit and Pindyck (1994). Time is divided into discrete intervals of length δ: accordingly, F ranges over a discrete set of values F_l, where l goes from 0 to u (which corresponds to the upper border of the continuation region). F_l makes discrete jumps $\varphi=F_{l+1}-F_l$.

First, suppose the boundary m and its discrete approximation u are given. Define

$$f = \left(f_l \middle| l = 0,....,u\right) \tag{14}$$

as the column vector of the squared forecast errors, plus adjustment costs (which arise at u). Then, let **L** be the column vector corresponding to the overall loss L, whose the component is L(F_l), the present value of losses plus control costs starting from state l.

$$\mathbf{L} = \sum_{q=0}^{\infty} f \exp(-\rho\delta q) \rightarrow \mathbf{L} = f + e^{-\rho\delta} L \tag{15}$$

From (15), it is easy to derive the value-matching condition, that holds at any boundary (for l=u), and states that the reduction in loss by exercising the control must match its cost.[7]

$$L(m) - L(0) = b \tag{16}$$

[7] See appendix 2.

I now turn to points inside the continuation region (i.e. for $l<u$, where no *thinking* costs are paid). In the relevant range of the continuation region, that is on the right-hand side of τ_2 in the figure above, the generic discrete F is

$$F_l = \frac{\sigma^2}{4\beta^2}\frac{1}{2k}\left[1 - e^{-2k(l-n)\delta} + \frac{e^{-2k(l-n)\delta}}{4}\left(1 - e^{-2kn\delta}\right)\right] \tag{17}$$

where $n\delta$ is the discrete-time equivalent of the difference τ_2-τ_1. Calculating the difference $\varphi=F_{l+1}-F_l$ and equating it to the discrete-time equivalent of the time differential of (12), it results that the following must hold:

$$F_{l+1} - F_l = \varphi = \frac{\sigma^2}{4\beta^2} e^{-2k(l-n)\delta}\left[1 - \frac{1}{4}\left(1 - e^{-2kn\delta}\right)\right]\delta \tag{18}$$

The Taylor series expansion of (18) around $\delta=0$ gives

$$\varphi = \frac{\sigma^2}{4\beta^2}\delta + o(\delta) \tag{19}$$

The same relationship (19) holds in the left-hand side of the continuation region (for points between τ_1 and τ_2).[8] Thanks to this result, the definition (15) can be used to show that for internal states the global loss obeys a first-order linear differential equation[9]

$$\frac{\sigma^2}{4\beta^2}L'(F) - \rho L(F) + F = 0 \tag{20}$$

The region over which the differential equation (20) holds is endogenous, since it depends on the boundary m, where the control is exercised (the producer *thinks*). Here, the so called smooth pasting condition must apply, that characterises the boundary as optimal by imposing the equality between the marginal cost of exercising and that of not exercising the control.[10]

[8] See Appendix 3.
[9] See Appendix 4.
[10] For a throughout theoretical discussion on the value matching and smooth pasting conditions in optimal stopping problems see Dixit and Pindyck (1993), the Appendix to Chapter 4.

$$L'\left(m^*\right)=L'(0) \tag{21}$$

where m* is the optimum boundary, that is the value taken by the expected squared forecast error F at time $t=\tau_1+\Delta$, being Δ the time lag between two successive applications of the control, i.e. two successive *thinking* times (see the figure). The solution of (20) consistent with (21) is

$$L(F)=\frac{\sigma^2}{4\beta^2\rho^2}+\frac{1}{\rho}F \tag{22}$$

The optimum boundary m* can now be derived. Write (22) for F=m* and F=0:

$$\begin{aligned} L\left(m^*\right)&=\frac{\sigma^2}{4\beta^2\rho^2}+\frac{1}{\rho}m^* \\ L(0)&=\frac{\sigma^2}{4\beta^2\rho^2} \end{aligned} \tag{23}$$

Substitute (23) in the value-matching condition (16), and the optimality condition obtains:

$$b=\frac{m^*}{\rho} \tag{24}$$

3.2 Oligopoly

The optimality condition obtained above is fairly general, since it holds regardless the number of firms. Indeed, the Taylor series approximation of $\varphi=F_{l+1}-F_l$ is $o(\delta)$ even if F is expressed for a generic number of firms as in (11), number that may range from 1 to infinity. In the following, I will analyse the equilibrium informational characteristics of the optimum in oligopoly, to be able to find an explicit expression for m*. The latter in fact turns out to be different depending on the producers' relative *thinking* times, which can in principle be the same for all (*thinking* contemporaneously), or not (*thinking* in groups and *thinking* at different times).

3.2.1 *Thinking* contemporaneously

The generic firm (let us say firm 1) must decide whether to *think* given that all the others are *thinking*. The relevant loss is then (11) with τ_j=t for j≠1 and τ_2-τ_1=t-τ_1=Δ. The optimality condition (24) becomes:

$$b = \frac{\sigma^2}{4\beta^2}\frac{1}{2k}\frac{1}{n^2\rho}(1-e^{-2k\Delta^*}) \tag{25}$$

Let us assume (25) is satisfied, so that all producers *think* contemporaneously. In such a case, however, the expected squared deviation of actual from full information output must be equal to (11) after posting equality of all the τ_i. On the boundary (at t=τ+Δ), equality (24) must hold, which implies:

$$b = \frac{\sigma^2}{4\beta^2}\frac{1}{2k}\frac{1}{\rho}(1-e^{-2k\Delta^*}) \tag{26}$$

Equations (25) and (26) are only consistent iff b=0, that is iff Δ*=0 and, trivially, in the case of monopoly (n=1). Hence, for positive *thinking* costs, this cannot be a Nash equilibrium.

3.2.2 *Thinking* in groups

There is the possibility that some firms *think* at the same dates, and the others not. I will make here the simplest case: all firms but one *think* at the same times.

Let us say that all the firms from 1 to n-1 last thought at the same τ_1, whilst firm n last thought at τ_n, with $\tau_1<\tau_n$. Now, let us analyse the choice of one of the n-1 firms, for example firm 1. at τ_1+Δ her expected squared deviation from full information output, given that the other firms are following the described strategy, would be equal to (11) with τ_j=t for j=2,…,n-1 and τ_2-τ_1=t-τ_1=Δ:

$$\begin{aligned} E[(q_1^*-q_1)^2] &= \frac{\sigma^2}{4\beta^2}\frac{1}{2k}\{1-e^{-2k(t-\tau_n)}+ \\ &+\frac{e^{-2k(t-\tau_n)}}{4}[1-e^{-2k(\tau_n-t)}]+\frac{1}{n^2}[1-e^{-2k\Delta}]\} = \\ &= \frac{\sigma^2}{4\beta^2}\frac{1}{2k}\{\frac{3}{4}[1-e^{-2k(\Delta+\tau_1-\tau_n)}]+\frac{1}{n^2}+[1-e^{-2k\Delta}]\} \end{aligned} \tag{27}$$

It would be optimal for firms from 1 to n-1 to update their information sets at $\tau_1+\Delta^*$ if and only if (27) verifies the optimality condition (24), that is if its right-hand side divided by the discount rate is equal to b. However, if all the firms except firm n *think* at the same times, their loss becomes:

$$E[(q_1^* - q_1)^2] = \frac{\sigma^2}{4\beta^2}\frac{1}{2k}\{1 - e^{-2k(t-\tau_n)} + \frac{e^{-2k(t-\tau_n)}}{4}[1 - e^{-2k(\tau_n-\tau_i)}]\} = \\ = \frac{\sigma^2}{4\beta^2}\frac{1}{2k}\frac{3}{4}[1 - e^{-2k(\Delta+\tau_1-\tau_n)}] \tag{28}$$

The optimality condition (24) requires that at $\tau_1+\Delta^*$ the right-hand side of (28) divided by the discount rate is equal to b. The two conditions can be jointly verified if and only if b=0, which would however imply that all the firms (n included) are *thinking* together all the time. Hence, *thinking* in groups cannot be an equilibrium.

3.2.3 *Thinking* at different times

Let us say that all producers *think* at different times, with a generic lag $a_i\Delta$ between τ_i and τ_{i+1}, where the sum of the a's is equal to 1.

```
   a1       a2      a3         .........            an   t
---+------+------+------+------+------+------+------+------+---
  τ1      τ2      τ3      τ4  .........          τn    τ1+Δ
```

After substituting the above notation into (11) and applying the optimality condition (24), the following results:

$$b = \frac{\sigma^2}{4\beta^2}\frac{1}{2k\rho}[1 - e^{-2ka_n\Delta^*} + \frac{e^{-2ka_n\Delta^*}}{4}(1 - e^{-2ka_{n-1}\Delta^*}) + \\ \frac{e^{-2k(a_n+a_{n-1})\Delta^*}}{9}(1 - e^{-2ka_{n-2}\Delta^*}) + \ldots + \frac{e^{-2k(a_n+\ldots a_2)\Delta^*}}{n^2}(1 - e^{-2ka_1\Delta^*})] \tag{29}$$

Condition (29) must be verified for all the firms once they approach their respective *thinking* times. i.e. it must apply with shifting a_i's. Hence, it must be true that the a_i's are all identical and equal to 1/n. Substituting into (29):

$$b = \frac{\sigma^2}{4\beta^2}\frac{1}{2k\rho}(1-e^{\frac{-2k\Delta^*}{n}})*$$
$$*[1+e^{\frac{-2k\Delta^*}{n}}(\frac{1}{4}+\frac{e^{\frac{-2k\Delta^*}{n}}}{9}+\frac{e^{\frac{-4k\Delta^*}{n}}}{16}+...+\frac{e^{-(n-2)\frac{2k\Delta^*}{n}}}{n^2})] \qquad (30)$$

The above expression can be also written as

$$b = \frac{\sigma^2}{4\beta^2}\frac{1}{2k\rho}(1-e^{\frac{-2k\Delta^*}{n}})\sum_{i=1}^{n}\frac{e^{(1-i)\frac{2k\Delta^*}{n}}}{i^2} \qquad (31)$$

Equation (31) identifies the only Nash equilibrium of this game. In equilibrium, every oligopolist *thinks* exactly at an intermediate point in time between the last *thinking* time of the previous and the next *thinking* time of the following producer. Hence, I shall call this a "*think* in the middle" strategy.

3.2.4 Intuition of the result and comparative statics

For positive finite values of the parameters on the right-hand side of (31), the optimal strategy is that of *thinking* continuously ($\Delta^*=0$) if and only if the cost b is equal to zero. In such a case, all the producers use the same information set all the time, that is they continuously take into account the true realization of the stochastic variable α when making decisions. Information is fully exploited and the solution maintains all the features of standard rationality.

If however b is positive, (31) expresses the optimal lag Δ^* as a function of the cost itself, the price elasticity β, the speed of convergence of the stochastic intercept towards its long-run mean k, the diffusion coefficient σ, the discount rate ρ and the number of firms n. In this case, a "*think* in the middle" strategy obtains, so that the producers will decide to use different information sets at any time. The intuition behind this result is that there is a positive externality of *thinking*, that however has nothing to do with information sharing. When a producer *thinks*, she gets a loss equal to b regardless what her opponents do. Instead, at times when she is not updating her information set, she is better off if the other oligopolists *think*. This happens because the agent who thinks acts as a sort of buffer, pushing the effective price towards the level expected by the others. This can give rise to a sort of "free riding" problem. A graphical example with only two firms can help to explain this idea.

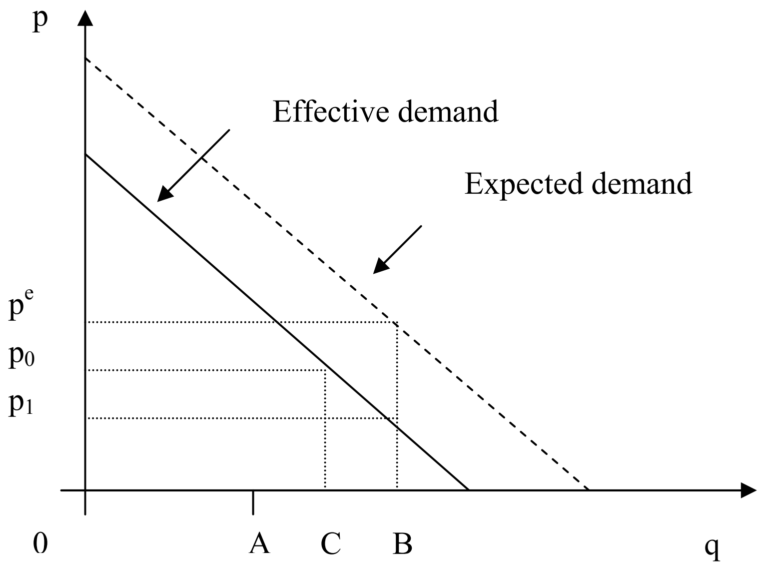

A negative demand shock pushes the demand function to its left. If the duopolist (let us say 1) does not *think*, she will not observe the shock and she will produce 0-A, i.e. more than would be optimal. Remember that 1's competitor can be equally or better informed. In either case, 1's best guess is that 2 produces as much as herself, (A-B)so that she expects the price to go to p^e. If 2 uses the same information set as 1 (*think* contemporaneously case), she will indeed produce A-B and the price would go to p_1. If on the contrary producer 2 *thinks*, she takes into account the negative demand shock and produces the smaller quantity A-C, so that the price goes to p_0, which is closer than p_1 to the price p^e, expected by producer 1. This also implies a smaller expected squared deviation of 1's production from its full information level, so that it implies a positive externality for the producer that does not *think*. In this sense, the *thinking* producer acts as a residual stabiliser, so that the marginal benefit of *thinking* for the other producer shrinks and, given the cost b, she prefers to wait. The "*thinking* in the middle" strategy is the only Nash equilibrium because it makes each producer able to exploit at best this externality.

Equation (31) can be used to perform some comparative statics. The effect of a larger cost b on the optimal thinking lag Δ^* is straightforward: as the cost of *thinking* rises, agents update their data sets less frequently. A rise in the instantaneous variance of the stochastic component α shortens the time lag between *thinking* times: more uncertainty is connected with bigger expected forecasts errors in absolute value, increasing the cost of inaction and making it preferable to update more frequently one's information set. The more reactive is the demand price to quantity (the bigger is β) and the less are people worried about the future (the higher is ρ), the wider is optimal Δ. Indeed, a larger elasticity raises Δ^* because a given forecast error gives rise to smaller losses: inappropriate output levels have smaller revenue consequences when demand is more elastic. The positive effect of a bigger discount rate on the optimal *thinking* lag is due to the fact that the producer attributes a larger weight to present *thinking* costs than to future forecast errors, and tends to postpone the cost itself. As far as the effect on equilibrium of the speed of convergence k is concerned, a faster speed of convergence (higher k) yields less frequent updating of the information set (bigger Δ^*): more rapid convergence implies smaller forecast errors and justifies longer periods of inaction. This effect however vanishes when $\Delta^*=0$, i.e. when agents continuously update their information sets.

Finally, since the right-hand side of (31) is decreasing with n and increasing with Δ, as the number of firms increases the optimal Δ must widen. As n tends to infinity, (as the oligopoly model tends to perfect competition), the marginal benefit of *thinking* tends to become negligible (the right-hand side of 31 goes to zero) so that the optimal *thinking* lag widens enormously. This is an interesting feature of the present model. It confirms Hwang's results, that competitive firms acquire less information than oligopolistic firms. Moreover, it points directly to Grossman and Stiglitz (1980) "impossibility of informationally efficient markets": if information collection or utilisation is costly, in a competitive environment nobody will have the incentive to acquire new information.

To get an intuition of the above result, it must be added that – as the number of firms increases - the optimal thinking lag Δ^* for any single firm rises, but the overall frequency of *thinking* is also higher. To make an example, a duopolist updates her information set more occasionally than a monopolist, but *thinking* by one of the two happens more frequently then *thinking* by the monopolist. This happens because global production in duopoly is bigger, so that on average it pays more to *think*.

As a consequence of the informational strategy adopted in equilibrium, the "buffer" effect exerted by the *thinking* firm, that pushes prices closer to the level expected by the other – *non-thinking* – producers, is enhanced when the number of firms increases. In a competitive environment (with many firms), almost at every time there is somebody who is *thinking*: this, in addiction to the fact that for given demand each firm can appropriate a smaller market share, renders the marginal benefit of *thinking* more and more negligible in comparison to its cost. As a consequence, the optimal lag between successive *thinking* times widens.

4. Conclusions

This paper extends the analysis and the methodology developed in a previous work to a model of oligopoly,[11] where agents must pay a fixed, sunk cost for using information about the state. This cost is referred to as the cost of *thinking,* intending the procedure through which people organise and then elaborate information. The optimisation renders a time-dependent rule, where *thinking* occurs at discrete time intervals: the optimal lag between each subsequent updating of the information set is increasing in the cost of *thinking* itself, while it is decreasing in the unconditional variance of the stochastic state variable (bigger uncertainty produces more intensive *thinking*).

The entirely new result here is that, even if the producers do not share information, an informational externality is present, so that firms will optimally *think* at different times: *thinking* is staggered, and asymmetric information becomes an endogenous feature of the equilibrium – consistently with Grossman and Stiglitz' (1980) outcome. It is important to note that the informational externality has nothing to do with the revelation of information through prices, since I assume that all information is virtually at the disposal of everybody at zero cost, and that the relevant cost must be paid for processing rather than for acquiring it.

I also find a negative correlation between the incentive for the firm to acquire information and the degree of competition on the market: monopoly and oligopoly firms use more information than competitive firms. Indeed, the amount of information that is gathered and used is endogenous, or better it is regime specific. Different market structures – and different forms of public interventions to combat market failures – will induce differences in the frequency of decision-making and the average quality of information on which actions are based. The design of policy must therefore have regard to its informational implications.

Isabella Imperato
Ministry of the economy and Finance – Rome

[11] Begg-Imperato (2001).

Appendix 1

In this demonstration, I will use the following property of the mean reverting process α

$$\alpha_t = E\left(\alpha_t \middle| \alpha_{\tau_i}\right) + \sigma e^{-k(t-\tau_i)} \int_{\tau_i}^{t} e^{k(s-\tau_i)} dz_s \tag{A1}$$

Applying (9) in the text to producer 1, the one who thought firs and is hence the less informed, I get

$$q_{1t} = \frac{1}{\beta(1+n)} E_1(\alpha_t) \tag{A2}$$

Using A2 in conjunction with equation (5) with 1 substituted for i

$$q_{1t}^* - q_{1t} = \frac{1}{2\beta}\left[\alpha_t - \beta \sum_{j=2}^{n} q_{jt}\right] - q_{1t} \tag{A3}$$

To compute the right hand side of A3, I first calculate the sum from firm 1 onwards, then subtract q_1. Applying recursively (9), the sum of production levels from 1 to a generic i is equal to

$$\begin{aligned}\sum_{j=1}^{i} q_{jt} = \frac{1}{\beta(n-i+2)}\{E_i(\alpha_t) + \frac{n-i+1}{n-i+3}[E_{i-1}(\alpha_t) + \frac{n-i+2}{n-i+4}[E_{i-2}(\alpha_t) + \\ + \frac{n-i+3}{n-i+5}[E_{i-3}(\alpha_t) + \ldots + \frac{n-2}{n}[E_2(\alpha_t) + \frac{n-1}{n+1}E_1(\alpha_t)]]]\ldots\}\end{aligned} \tag{A4}$$

The sum from 1 to n can be found substituting i with n, which gives

$$\sum_{j=1}^{n} q_{jt} = \frac{1}{2\beta}\{E_n(\alpha_t) + \frac{1}{3}[E_{n-1}(\alpha_t) + \frac{2}{4}[E_{n-2}(\alpha_t) + \frac{3}{5}[E_{n-3}(\alpha_t) + \ldots + \frac{n-2}{n}[E_2(\alpha_t) + \frac{n-1}{n+1}E_1(\alpha_t)]]]\ldots\} = = \frac{1}{2\beta}[E_n(\alpha_t) + \frac{1}{3}E_{n-1}(\alpha_t) + \frac{1}{3}\frac{2}{4}E_{n-2}(\alpha_t) + + \frac{1}{3}\frac{2}{4}\frac{3}{5}E_{n-3}(\alpha_t) + \ldots + \frac{1}{3}\frac{2}{4}\frac{3}{5}\ldots\frac{n-2}{n}\frac{n-1}{n+1}E_1(\alpha_t) \tag{A5}$$

Subtracting A2 from A5 and simplifying I get

$$\sum_{j=2}^{n} q_{jt} = \frac{1}{2\beta}[E_n(\alpha_t) + \frac{1}{3}E_{n-1}(\alpha_t) + \frac{1}{3}\frac{2}{4}E_{n-2}(\alpha_t) + \ldots \ldots + \frac{2}{n(n-1)}E_2(\alpha_t) - \frac{2}{1+n}\frac{n-1}{n}E_1(\alpha_t)] \tag{A6}$$

and, finally, substituting A6 and A2 into A3:

$$q_{1t}^{*} - q_{1t} = \frac{1}{2\beta}\{\alpha_t - E_n(\alpha_t) + \frac{1}{2}[E_n(\alpha_t) - E_{n-1}(\alpha_t)] + \frac{1}{3}[E_{n-1}(\alpha_t) - E_{n-2}(\alpha_t)] + \ldots + \frac{1}{n}[E_2(\alpha_t) - E_1(\alpha_t)]\} \tag{A7}$$

Using A1, the right-hand side of A7 becomes

$$\frac{\sigma}{2\beta}e^{-kt}\left[\int_{\tau_n}^{t} e^{ks}dz_s + \frac{1}{2}\int_{\tau_{n-1}}^{\tau_n} e^{ks}dz_s + \frac{1}{3}\int_{\tau_{n-2}}^{\tau_{n-1}} e^{ks}dz_s + \ldots + \frac{1}{n}\int_{\tau_1}^{\tau_2} e^{ks}dz_s\right] \tag{A8}$$

whose expected squared value is equal to (11) in the text.

Appendix 2

Definition (14) in the text implies that, for the extreme state l=u, $f_l=b+\delta F_l$, so that (13) becomes

$$L(F_l) = b + F_l\delta + \exp(-\rho\delta)L(F_{l+1}) \quad \text{(A9)}$$

Multiplying by exp(ρδ) and rearranging terms we get

$$\exp(\rho\delta)L(F_l) - L(F_{l+1}) = \delta\exp(\rho\delta)F_l + \exp(\rho\delta)b \quad \text{(A10)}$$

Where $\varphi=F_{l+1}-F_l$ is o(δ).[12] Equation (A10) can be expanded using a Taylor series. Its left-hand side expands to

$$[1+\rho\delta+o(\delta)][L(F_l)+L'(F_l)\varphi+o(\varphi)] - L(F_{l+1}) \quad \text{(A11)}$$

And its right-hand side expands to

$$\delta[1+\rho\delta+o(\delta)]F_l + b[1+\rho\delta+o(\delta)] \quad \text{(A12)}$$

Equation (A11) can be written as

$$L(F_l) - L(F_{l+1}) + \Omega(\delta) \quad \text{(A13)}$$

Where Ω(δ) comprises terms such that Ω(δ)/δ stays bounded above as δ approaches zero. Using this notation, (A12) becomes

$$b + \Omega(\delta) \quad \text{(A14)}$$

Equating (A13) to (A14) and letting δ go to zero I get the value matching condition (16) in the text.

[12] See equation (19) in the text

Appendix 3

Between τ_1 and τ_2, the squared deviation of production from its full information level F must be computed under the assumption that producer 1 is more informed than producer 2. Using (5), after substituting to α its expected value, her effective production is equal to

$$q_{t1} = \frac{1}{2\beta}\left[E\left(\alpha_t \middle| \alpha_{\tau 1}\right) - \beta q_{t2}\right] \tag{A15}$$

Her full information production can be easily derived form (5), so that the squared divergence F becomes

$$E\left[\left(q_{t1}^{*} - q_{t1}\right)^2\right] = \frac{1}{4\beta^2}\frac{\sigma^2}{2k}\left[1 - e^{-2kt\left(t-\tau_1\right)}\right] \tag{A16}$$

Applying to (A16) the same discretisation process used for (17), (19) obtains.

Appendix 4

For internal states (l≤u-1), the global loss (15) becomes

$$L\left(F_l\right) = F_l\delta + \exp\left(-\rho\delta\right)L\left(F_{l+1}\right) \tag{A17}$$

Multiplying both terms by exp(ρδ) and rearranging, (A17) becomes

$$\exp\left(\rho\delta\right)L\left(F_l\right) - L\left(F_{l+1}\right) = \exp\left(\rho\delta\right)F_l\delta \tag{A18}$$

whose right-hand side expands to

$$\delta F_l + o\left(\delta\right) \tag{A19}$$

The left-hand side can be re-written as

$$L(F_l)[\exp(\rho\delta)-1]-[L(F_{l+1})-L(F_l)] \qquad \text{(A20)}$$

whose first-order expansion is

$$\rho\delta\, L(F_l)-\left[L^{'}(F_l)\varphi+o(\varphi)\right] \qquad \text{(A21)}$$

Now, put (A19) and (A21) together, substitute φ with (19), divide both terms by δ and let δ go to zero. Replacing F_l by F, the present value of losses obeys the first-order linear differential equation (20) in the text.

References

Begg, D.K.H. (1982). *The rational expectations revolution in macroeconomics. Theories and evidence*, Hemel Hempstead, Philip Allan.

Begg D.K.H., Imperato I. (2001), "The rationality of information gathering: monopoly", **The Manchester School**, No.69, pp.237-252.

Board, R. (1994). "Polynomially bounded rationality", **Journal of Economic Theory,** No. 43, pp. 246-270.

Bray, M. (1982). "Learning, estimation, and the stability of rational expectations equilibria", **Journal of Economic Theory**, No. 26, pp. 318-339.

Caballero, R.J. (1989). "Time dependent rules, aggregate stickiness and information externalities", **Columbia University Discussion Paper Series**, No. 428.

Crettez, B., Michel, P. (1992). "Economically rational expectations equilibrium", **Economic Letters**, No. 40, pp. 203-206.

Dixit, A. (1991). "A simplified treatment of the theory of optimal regulation of Brownian motion", **Journal of Economic Dynamics and Control,** No. 15, pp. 657-673.

Dixit, A. (1993). “The art of smooth pasting”, in Lesourne J. and Sonnenschein H. (Eds.), *Fundamentals of pure and applied economics,* Chur, Harwood Academic Publisher, Vol. 55.

Dixit, A., Pindyck, D.K. (1994). *Investment under uncertainty,* Princeton, Princeton University Press.

Dwyer, G.P., Williams, A.W., Battalio, R.C., Mason, T.I. (1993). “Tests of rational expectations in a stark setting”, **Economic Journal,** No.103, pp. 586-601.

Evans, G.W. (1985). “Expectational stability and the multiple equilibrium problem in linear rational expectations models”, **Quarterly Journal of Economics**, No.100, pp. 1217-1233.

Evans, G.W., Honkapohja, S. (1994a). “Learning, convergence and stability with multiple rational expectations equilibria”, **European Economic Review**, No. 38, pp. 1071-1098.

Evans, G.W., Honkapohja, S. (1994b). “On local stability of sunspot equilibria under adaptive learning rules”, **Journal of Economic Theory**, No. 64, pp. 142-161.

Evans, G.W., Honkapohja, S. (1995). “Adaptive learning and expectational stability: an introduction”, In Kirman, A., Salmon, M. (Eds.), *Learning and rationality in economics*, Oxford, Blackwell, pp. 102-126.

Evans, G.W., Ramey, G. (1992). “Expectation calculation and macroeconomic dynamics”, **American Economic Review,** Vol. 82, No. 1, pp. 207-224.

Feige, E.L., Pearce, D.K. (1976). “Economically rational expectations: are innovations in the rate of inflation independent of innovations in measures of monetary and fiscal policy?”, **Journal of Political Economy**, No. 84, pp. 499-522.

Fleming, W.H., Rishel, R.W. (1975). *Deterministic and stochastic optimal control*, New York, Springer-Verlag.

Fourgeaud, C., Gourieroux, C., Pradel, J. (1986). “Learning procedures and convergence to rationality”, **Econometrica**, Vol. 54, No. 4, pp. 845-868.

Fuhrer, J.C. (1987). “Information gathering and expectation formation under model uncertainty”, **Southern Economic Journal**, No. 5, pp. 685-701.

Grossman, S.J., Stiglitz, J.E. (1980). “On the impossibility of informationally efficient markets”, **American Economic Review**, Vol. 70, No. 3, pp. 393-408.

Kirman, A., Salmon, M. (Eds.) (1995). *Learning and rationality in economics*, Oxford, Blackwell.

Krylov, N.V. (1980). *Controlled diffusion processes*, New York and Berlin, Springer-Verlag.

Mankiw, N.G., Romer, D. (Eds.) (1991). *New Keynesian Economics. Imperfect competition and sticky prices*, Cambridge MA, MIT Press.

Marcet, A., Sargent, T.J. (1988). “The fate of systems with adaptive expectations”, **American Economic Review**, Vol. 78, Papers and Proceedings, pp. 168-172.

Marcet, A., Sargent, T.J. (1989). “Convergence of least squares learning mechanisms in self-referential stochastic models”, **Journal of Economic Theory**, No. 48, pp. 337-368.

Muth, J.F. (1961). “Rational expectations and the theory of price movements”, **Econometrica** Vol. 29, No. 3, pp. 315-33.

Timmermann, A. (1994). “Can agents learn to form rational expectations? Some results on convergence and stability of learning in the UK Stock Market”, **Economic Journal,** Vol. 104, No. 425, pp. 777-797.