The emergence and survival of rational expectations

Temporary- do not quote

Marco J. Lombardi
European Central Bank

Luca Onorante[♣]
European Central Bank

Abstract

In an economy which admits unique rational expectations (RE) equilibrium, we study the effect of the insertion of bounded rational agents on the equilibrium of the economy and on the payoffs of the rational agents. Starting from RE, we investigate the interaction of rational and non rational agents when the latter are introduced. Then, starting from a non-rational random state, we explore the emergence of RE in the model. We show that, depending on the deep parameters and in particular on the degree of forward lookingness of the model, agents will converge quickly around the unique RE equilibrium or spend there most of the time with occasional drifts. Finally, we study the dynamics of the system "out of equilibrium" when the model is subject to periodical structural changes in the deep parameters.

1. Introduction

The recent financial crisis and the subsequent global economic downturn has led many commentators to question the foundations of standard macroeconomic models, and its usefulness in explaining real-world developments as the one we have just experienced. In particular, the realism of the idea that agents of the economy are homogenous and behave rationally has been questioned as unrealistic and unable to explain the gyrations and the formation of bubbles in asset prices. <NEED THIS?>

[♣] Corresponding author. E-mail: luca.onorante@ecb.europa.eu. The views expressed are solely those of the authors and do not necessarily reflect those of the European Central Bank or the ESCB.

In this paper, we try and explain the reasons of the emergence of occasional strong macroeconomic fluctuations by assuming that the economic agents can form heterogeneous expectations according to different models. We do this by modeling the the agents within a genetic algorithm, which is capable of producing both RE and completely ad-hoc expectations. We proceed as follows. First, we explain our choice of a GA by comparing it with other possible relaxation of RE. We will describe the GA as a technique of loosening the strict RE paradigma while retaining the attractive feature that the agents are able to optimize given their beliefs, and argue that a GA has some advantages over other similar techniques because it is less demaning in terms of hypotheses, it fits more data (in particular, expectations) and more accurately, it can be easily mapped into activities observed at the micro level, and it is microfounded in a loose sense.

In the second part of the paper, after having introduced and justified the use of the GA, we describe its actual functioning and propose our results.

The notable results are the following: persistence in the data series is generated endogenously, without recurring to autocorrelated disturbances; heterogeneous expectations are generated and compared with existing expectations surveys; finally, and most importantly, we are able to explain under which conditions occasional strong deviations from equilibrium may arise and "surprise" the markets and the forecasters. The use of a GA also provides with interesting insights when the models are subject to structural changes. These will be the topic of further research.

2. Relaxing the rational expectations: a non exhaustive review

In a classic model with rational expectations (RE), agents must:

1. be perfectly informed about the model governing the economy, including preferences, technology and any structural parameters;
2. they must of course behave rationally;
3. and finally they need to be able to coordinate with other agents.

The latter requirement implies that they have common certainty of rationality, or that they are convinced that all other agents are rational. Under these requirements a representative agent hypothesis can be used to solve the model. These hypotheses are widely used in macroeconomic modelling. This comes despite their restrictiveness, the fact that they are at odds with a growing stock of evidence provided by experimental economics and against the background of the difficulty of the models based on RE to explain some relevant features of macro data.

Among the problems with data matching, well-known problems of RE are that they a) are unable to endogenously generate persistence in the economic series, b) lead to counterintuitive conclusions as to how some variables (e.g. consumption) react to external shocks, c) do not reproduce other relevant features such as the difference in length between expansions and recessions. Most RE models also d) do not produce economic fluctuations beyond those implied by the exogenous shocks.

A nowadays standard way to enrich RE models leading them towards more realism is the introduction of missing markets, frictions, outocorrelated shocks or other forms of persistence. From a theoretical standpoint point of view, assumptions 1 to 3 above remain valid, but the model is enriched with:

4. persistence (“frictions”) coming from the outside of the economy
5. the assumptions that the agents are aware and take into account this persistence when deciding.

This way to overcome the limits of RE has proven very successful in the theoretical literature, as the introduction of restrictions allows to better explain and describe relevant economic mechanism without modifying the underlying RE implant. Yet, when it comes to describe relevant features of the data, it turns out that the amount of necessary frictions becomes quite considerable, making the overall number of hypotheses of the kind described in points 4 and 5 very high (quote a model which has gone a long way towards the data, SW, needs 4 different frictions plus habit persistence MAYBE ALSO DESCRIBE TYPES OF FRICTIONS AND THEIR REASONABLENESS). Finally, another unconvincing feature of the use of frictions is that the persistence is postulated to be exogenous from the model, and is not endogenously created.

An alternative approach constrains the capacity of (some of) the economic agents to decide optimally. For examplea common modelling solution (first introduced by Gali et a., XXXX) divides the agents of the economy between RE agents and constrained agents, who are e.g. “forced” to spend entirely their current income, or are more generally acting in some other suboptimal way. Other models QUOTE assume that agents are limited in their capacity to optimize by habit persistence. As before, these solutions also do not imply any relaxation of the hypotheses 1 to 3, as at least some of the agents are able to find the RE equilibrium. Instead, the additional hypothesis are that some of the agents are constrained in their capacity to decide or to act. Yet, as before, it is postulated that the rational agents are able to find the RE equilibrium even in presence of "dumb" agents, whose presence is hence taken into account in their optimization strategy. In order to be able to do that, rational agents they must not only know the model, but also the number and the behaviour of all the (rational and non-rational) agents, which constitutes an enormous amount of information to be taken into account when they solve the model.

The common feature of both approaches outlined above is that they maintain the RE hypothesis, and add further assumptions for the sake of realism. Deciding between them is in the end a matter of comparing the number and restrictiveness of the hypotheses on one side and the additional insights and realism on the other side.

A different approach has focused on loosening hypotheses 1 and 3. In particular, criticism about the realism of hypothesis 1, which implies that the agents in the model actually have much more information than the economist who created it and estimated it on real data (quote the conquest of american inflation Sargent), along with the attempt of reproducing realistic features of the data, leads to the substitution of this assumption with the hypothesis that agents do not have perfect information, but rather they may learn and they behave like econometricians, or Bayesian learners. Most of the contributions, for example Townsend (1978, 1983), Bray and Kreps (1986), Brandenburger (1984a) still assume that the agents know (at least roughly) the structure of the model and are simply learning about the parameters as they observe data, but this is not a necessary feature and there are also studies on the effect of learning with misspecified models. For instance, Sims (71) and White (1982, 1994) study the behaviour of agents behaving like econometricians with misspecified models in a context where the expectations do not influence the true model[1]. In Sargent (1999) and more recently Adam (XXXX), expectations determine the structure of the model itself.

Another positive feature is that hypothesis 2 (i.e. rationality given the beliefs) is maintained.

[1] In other contributions the agents learn using specifications of the model that are either correct only in the RE equilibrium and incorrect in the process of convergence (Blume and Easley (1982), Bray (1982, 1983), Radner (1982), Frydman (1982), Fourgeaud, Gourieroux, and Pradel (1986)), Thisb biblio from Bray Savin 1986, read papers). The requirement imposed on the agents' knowledge are weaker, but sometimes the limitations are unnatural.

Relaxing hypothesis 1 comes at the expense of a certain degree of arbitrariness, as the agents may now be assumed to use several different heuristics to gather information, and only realism can distinguish among them. On the positive side, the hypothesis of learning replaces hypothesis 1 instead of adding on top of it.

In a nutshell, such modelizations of the agents have the advantage of reproducing a highly realistic feature of economics, i.e. that parameters of model have to be estimated, while keeping some discipline: the agent and the econometricians use (or at least they can use) the same amount of information, as they have both agreed upon a certain structure of the model, and update their beliefs about parameters using the same dataset. From an empirical standpoint, the interesting feature of this approach is that it can generate very rich dynamics. Both these characteristics put them into an advantage compared to models using frictions or constrained agents.

3. Relaxing RE using Genetic Algorithms

The modelling device we choose is a Genetic Algorithm (henceforth GA). A genetic algorithm is (along with evolution strategy and genetic programming) a class of stochastic search algorithms based on analogies with biological evolution.

In this section we consider genetic algorithms as one of different possible loosening of the hypotheses underlying RE and we compare it with other modifications of the pure RE paradigm, namely the introduction of agents facing constraints, the introduction of frictions, the use of learning.

As a relaxation of RE, a GA model has merits which are similar to those of learning: it eliminates some very restrictive hypotheses. The difference (and possibly the advantage of this modeling solution) is that the hypothesis 1 of knowledge of the model can be not only relaxed but even completely removed, as agents may be assumed to ignore the structure of the model. In a GA model, even in those extreme versions where the agents do not know the correct model of the economy, they are able use the GA as a heuristic to gather information.

Hypothesis 2 of rationality is maintained, in the sense that agents are still able to optimize on the basis of their beliefs. The main disciplinary device is kept, in that the actions of the agents in the model are optimal given their beliefs.

Given that in a GA the agents have almost by definition heterogeneous beliefs, the hypothesis 3 of coordination of the expectations is also not necessary, as the representative agent is abandoned. This generalization comes with a price: GA models cannot be analytically solved but only simulated. We do not believe this is a real problem for several reasons: first, most RE models (with the exception of the most simple ones) are also solved using simulation techniques; second, modern computers have made these techniques inexpensive; third and most important, the GA is a well known technique used in several fields such as biology and engineering which relies on a solid theoretical background and provides reliable results.

What are the modelling advantages of a GA? Many advantages are in common with learning. Both GA and learning produce persistence in variables without recourse to exogenous hypotheses (such as autocorrelated shocks). GA operators (imitation, mating etc) correspond to know heuristics described in micro theory just like learning corresponds to known ways of getting to know the parameter of the data. Both approaches are less restrictive than RE because they do not suppose that the agent knows more than the modeller, but GA can go as far as allowing the agents to ignore important features (such as the structure) of the model. Quite importantly to explain periods of turbulence, GAs are able to produce endogenous fluctuations beyond those imposed with the shocks, propagating them not only in time but also increasing their size. A final important advantage of GAs is that they provide a technology which makes it easier and almost natural to model heterogeneous agents. Thus, uncertainty as disagreement and biases in expectations can be part of the model. This is a relevant aspect in empirical assessment of a model, since expectations and macroeconomic variables are modelled together and the expectations produced by the model can be compared with those measured from surveys.

The model

We use the simplest possible framework, an economy expressed by one equation only (a “Phillips curve”):

$$p(t) = a + b * E(p(t)) + u(t)$$

We therefore abstract from any pretension of realism, in favour of a model in which the outcomes can be clearly attributed to the action of the agents rather than to the equations of the model. In doing this, we follow Sargent (1999), who uses as a workhorse a version of Bray's (1982) pricing model. As in Sargent, we emphasize a relevant characteristic that will be the object of our study, namely the forward-looking term in the equation. This feature is essential: without a forward looking term the economy would not be influenced at all by the expectations of the agents, and rational, bounded rational and even purely adaptive agents would simply converge to the fixed equilibrium at different speeds. The true model (the actual law of motion) varies depends on how the agents' model is specified. We also allow for backward looking expectations in some versions of our model.

We assume that the agents use a structural form of the model compatible with rational expectations, but they do not know the parameters. To simplify the interpretation of the results, and without loss of generality, all the agents have the same preferences, and differ only by how they form their expectations.

The expectations of the agents evolve within the GA. Within the GA, agents can also choose to be true RE-believers. The RE-believers are given the correct model of the economy and immediately reach the RE equilibrium. However, when other types of agents are introduced, RE-believers erroneously continue to assume that all the other agents will behave as they do[2], while this is not true because agents who do not know the true model of the economy search for the best possible prediction by using heuristic methods.

Each agent can also choose to use another model produced in the GA. Reasons to use a model other than RE include 1) ignorance of the exact parameters of the model, 2) lack of knowledge about the preferences of the other agents, or 3) awareness that other agents have heterogeneous preferences. If either of these conditions is met, the agent (knows that he) does not know the true model of the economy; furthermore, those who believe that expectations play some role in determining outcomes must also be aware that the optimum is a shifting one.

For analytical purposes, while running our GA, at each point in time we keep track of three 'types' of agents: 1) those choosing rational expectations, 2) those using other models, and within the latter 3) those ending up (via imitation and mutation) having a model to form expectations which is formally identical as the RE. We call the second group non-RE and the last group RE-imitators. It turns out that the RE-imitators are essential in speeding up the process of equilibrium when they show a

[2] The alternative hypothesis would be a very demanding one, as discussed in the previous section.

‘herd behaviour’ towards RE, but may also be responsible of large and protracted deviations from it.

This model departs from the existing literature in three dimensions. First, it produces endogenously varying shares of both “rational” and “non-rational” (heterogeneous) agents. As a matter of fact, in our model a high and varying heterogeneity of agents is not only permitted, but is an essential ingredient. It also leads to empirically testable implications about first and second moments of expectations.

Second, the evolution of the model based on a GA describes in a stylized way well-known micro phenomena of learning such as survival of the best ideas over time, modification of existing ones, experimentation, and imitation.

Third, we differentiate between genotype (the model used internally by the agent) and phenotype (the forecast apparent to the others). Agents can get closer to the right forecast either by using RE or by learning, thus making the same prediction in equilibrium and obtaining the same payoff, but in fact using completely different models.

3.1 The implementation of the Genetic algorithm

A typical genetic algorithm is linked to the problem it tries to solve by two elements: a genetic representation of the solution domain, encoded in “chromosomes”, and a fitness function to evaluate the solution domain, called evaluation function.

Our chromosomes are represented by a mixture of real and binary values[3]. The structure of a chromosome i in our exercise is the following one:

xi = [Slopei(real)   Intercepti(real)   UsingREi(binary)],

or in compact notation:

xi = [Si, Ki, REi]

The chromosome x is used to form expectations at each time t as follows:

Ex(pt+1) = Si * pt + Ki   iff REi = false

= Rational expectation   iff REi = true

[3] In the traditional encoding (Holland XXX) the chromosomes or tentative solutions are represented by strings of ones and zeros. However, many different coding techniques exist (Davis 1991 XXXX).

It should be noted that RE-like expectations can arise given an appropriate parametrization even when REi is false: we call the agents with these expectations the RE-imitators. For RE-believers, REi is true and the parameters will be ignored in favour of RE . Ex(pt+1) is the expectation of p at time t+1 performed at time t using the parameters of the chromosome. The evaluation function simply assumes that the fitness (payoff) of a chromosome is higher when the forecast is accurate:

ft(xi) = -[Exi(pt) – pt]^2

where pt+1 is the realization of the variable. This general formulation is for explanatory purposes, and can be substituted with a pay-off based on a welfare function when using a specific model.

<BOX BEGINS HERE?>

The operators driving the GA are defined as follows:

- **Selection** replicates the chromosome into the next generation, or kills it, the probability of the first outcome increasing with fitness ft(xi)

  SEL(xi,t) = {Xi,t+1, 0}

  Alternatively, we tried selection as a tournament in which the fittest survives:

  SEL(xi,t, xj,t) = {xi,t+1} iff ft(xi) => ft(xj)
  = {xj,t+1} iff ft(xi) < ft(xj)

- **Mating** takes at random values from one or the other:

  MAT(xi,t,xj,t) ={[Si + I *(Sj – Si), Ki+I *(Kj–Ki),REi+ I *(REj – REi)]},

  where each I is an indicator function taking value 1 or 0

- Finally, the **mutation** operator changes slightly the paprmetes, or flips the 1/0 flag for rational expectations:

  MUT(xi)= {[Si + I *e1, Ki + I *e2, I* REi + I *((1-REi) – REi)]}

  where each I is an indicator function taking value 1 or 0.

<BOX ENDS HERE?>

The genetic algorithm proceeds in the following steps:

1. Choose the number of chromosomes which in each generation will survive selection (Ns), be product of mating (Nm), and a probability of mutation (Pm).
2. An initial population of chromoseomes of size N = Ns + Nm is generated at time 0.

x1, 0, x2, 0, … xN, 0

3. The fitness of each individual chromosome is calculated
   f0(x1,0), f0(x2,0), …. f0(xN,0)
4. Select chromosomes for selection from the current population. Highly fit chromosomes have a higher probability of surviving into the next generation.
5. Repeat the previous step until Ns chromosomes have passed selection.
6. Select couples of chromosomes for mating from the current population. Highly fit chromosomes have a higher probability of being selected. The offsprings are assigned to the next generation.
7. Repeat the previous step until Nm offsprings have been generated.
8. Apply to each gene of the new population the mutation operator with probability Pm.
9. The newly formed population becomes the current one, and steps from 3 onwards are repeated.

4. Simulation experiments

In this section we will present our main findings coming out of detailed simulation exercises. All experiments will be based on a set of 1000 Monte Carlo replications, each featuring 1000 agents, for time series of length 1000. We start from the RE equilibrium, and discard the first 100 periods.

As anticipated, for each point in time we keep track of the three ‘types’ of agents, their parameters, and the equilibrium they produce.

4.2 Endogenous cycles

One interesting feature of the way agents interact is that they are able to generate endogenous cycles. What happens indeed is that, when rational and pseudo-rational agents dominate the population, there is no longer ay comparative advantage in being orthodoxically rational, as the equilibrium is not affected. Therefore, rational agents start to decrease, and the presence of more irrational agents shifts the equilibrium. When the shift is significant, it becomes again convenient to be rational, and rational agents again prevail.

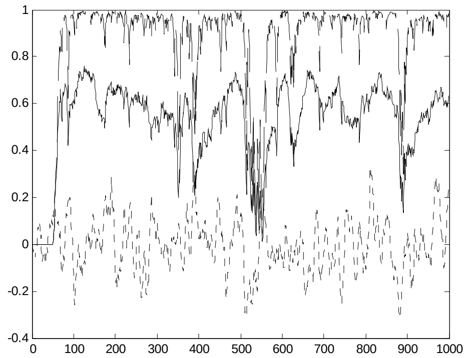

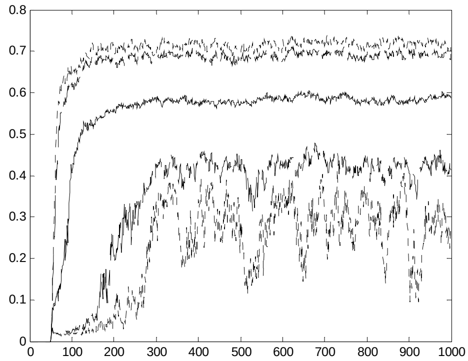

Figure 3. Example of endogenous cycle: equilibrium (dotted line) and percentage of rational agents (solid line).

Figure 4. Median and Quantiles of Monte Carlo replications.

QQQ

4.3 Endogenous persistence

How it arises when the model is only fwd looking
How small exo persistence is multiplied

4.4 How a crisis arises

Number of RE, var(expectations), var(models), var shock and effect of the shock, IRF depending on the initial conditions at shock. They are bigger than the initial shock.

4.5 How does the model compare with the data

5. Conclusions (see whether they fit)

While running our GA, at each point in time we keep track of three ‘types’ of agents: those using rational expectations, those using the GA, those using the GA but ending up (via imitation and mutation) having the same model to form expectations as the RE. It turns out that this subgroup of GA agents, showing a ‘herd behaviour’ towards RE, are essential in speeding up the process of equilibrium but are also responsible of large and protracted deviations from it.

Corresponding to our expectations, the (unique) RE equilibrium emerges for most parameterizations of the model, but the convergence takes more time in forward-looking models. We observe that in most cases[4] rational agents have a higher fitness when the system is out of equilibrium; therefore other agents either switch to true RE or they simply adjust the parameters in their models to produce expectations which are close to the RE.

We also observe that, once attained, the RE equilibrium can be sustained for protracted periods despite the presence of a small number of rational agents. To understand this point, it is important to note that when the RE equilibrium is attained, all non rational agents have also RE-compatible expectations, despite not using the RE model. Therefore, as long as the effect of mutation is negligible, the choices of RE and boundedly rational agents are observationally equivalent: RE agents receive the same fitness than those who are just imitating their choices and their number in the total population becomes a random walk. The RE equilibrium may hold for protracted periods even when all the RE agents have disappeared.

However, over time the effects of mutation become important and move the equilibrium away. Depending on the parameterization of the model (more precisely on the degree of forward-lookingness of the model) and on the degree of coordination of the expectations of the non rational agents, the system can drift away for quite prolonged amounts of time. These periodical deviations cannot be permanently sustained, as the potential advantage of RE agents in term of fitness increases with the distance of the system from the equilibrium.

Following a structural break, the RE agents instantaneously adjust, while the boundedly rational agents cannot do it. The new dynamic to equilibrium sees an increasing number of agents either adopting RE or imitating them. The variance in the expectations increases in the transition phase, in accordance with the empirical observation.

[4] Sargent (1999, the conquest of American inflation) shows that if agents behave like econometricians and use a fixed gain algorithm the RE equilibrium is substituted by a stationary stochastic process. The true model (the actual law of motion) varies depending on how the agents' model is misspecified. Our model reproduces this result, but due to the presence of heterogeneous expectations we show that these equilibria are temporary and relatively infrequent.