Where Do Booms and Busts Come From?[1]

Paul De Grauwe[2]

Prisme N° 22

October 2011

[1] I am grateful to Yunus Aksoy, Tony Atkinson, Robert Boyer, Casper de Vries, Stephan Fahr, Daniel Gros, Richard Harrison, Romain Houssa, Gerhard Illing, Pablo Rovira Kaltwasser, Christian Keuschnigg, Alan Kirman, Giovanni Lombardo, Lars Ljungqvist, Patrick Minford, John Muellbauer, Ilbas Pelin, Frank Smets, Robert Solow, Leopold von Thadden, David Vines, Volker Wieland and Tony Yates for their comments and suggestions.

[2] Paul De Grauwe is Professor of International Economics at the University of Leuven, Belgium. He is director of the money, macro and international finance research network of the CESifo (Munich Society for the Promotion of Economic Research) at the University of Munich, and a research fellow at the Centre for European Policy Studies in Brussels.

Summary

Capitalism is characterized by booms and busts. Periods of strong growth in output alternate with periods of declines in economic growth. Every macroeconomic theory should attempt to explain these endemic business cycle movements. In this text, I present two paradigms that attempt to explain these booms and busts. One is the Dynamic Stochastic General Equilibrium (DSGE) paradigm, in which agents have unlimited cognitive abilities. The other paradigm is a behavioural one, in which agents are assumed to have limited cognitive abilities. These two types of models produce radically different macroeconomic dynamics. I analyse these differences. I also study the different policy implications of these two paradigms.

1. Introduction

Capitalism is characterized by booms and busts; by periods of strong growth in output followed by periods of declines in economic growth. Every macroeconomic theory should attempt to explain these endemic business cycle movements.

The mainstream macroeconomic models, based on the supreme rationality of agents and their capacity to fully understand the complexity of the world, have told us that booms and busts are the result of exogenous shocks coming from outside the realm of the macroeconomy. Rationality prevails and the latter is peaceful. Unfortunately, outside shocks perturb the peaceful optimizing behaviour of rational agents and force them to change their plans. Thus, the mainstream explanation of booms and bust can be called a "meteor theory" of the business cycle. Once in a while a "big meteor" hits the macroeconomy leading to sudden disruption of economic activities.

This is not a very satisfactory theory. It shifts the burden of explaining why booms and busts in economic activity occur to other sciences. Macroeconomists should have the ambition of explaining booms and busts endogenously. They should analyse the dynamics that are present in the macroeconomy and that can lead to booms and busts endogenously. That is what I attempt to do in this article. I will present a behavioural macroeconomic model in which "Animal Spirits", as defined by John Maynard Keynes,[3] take centre stage.

Before developing the model, it is useful to present some stylized facts about the cyclical movements of output. Figure 1 shows the movements of the output gap[4] in the USA since 1960. We observe strong cyclical movements. They imply that there is strong autocorrelation in the output gap numbers, that is, the output gap in period t is strongly correlated with the output gap in period $t-1$. The intuition is that if there are cyclical movements, we will observe clustering of good and bad times. A positive (negative) output gap is likely to be followed by a positive (negative) output gap in the next period. That is what we find for the US output gap over the period 1960–2009: the autocorrelation coefficient is 0.94. Similar autocorrelation coefficients are found in other countries.

---

[3] Editor's note: In his 1936 book, *The General Theory of Employment, Interest and Money*, Keynes uses the term "animal spirits" to describe human emotion that drives investor confidence.

[4] The output gap refers to the difference between actual and potential gross domestic product (GDP).

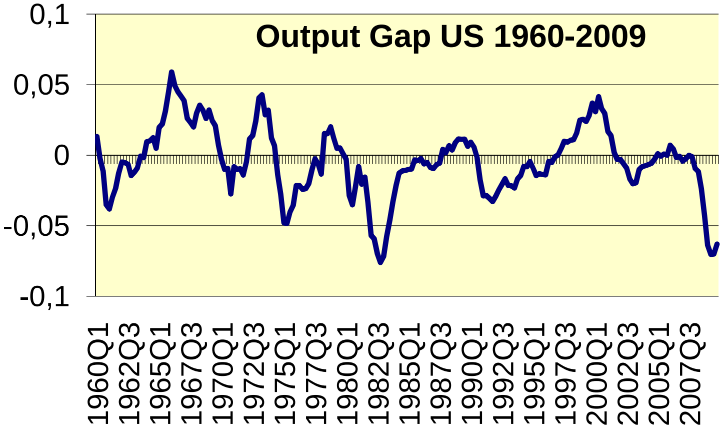

Source: US Department of Commerce and Congressional Budget Office

Figure 1: Output gap for the USA, 1960–2009

A second stylized fact about the movements in the output gap is that these are not normally distributed. The evidence for the USA is presented in Figure 2. We find, first, that there is excess kurtosis (kurtosis $= 3.62$), which means that there is too much concentration of observations around the mean to be consistent with a normal distribution. Second, we find that there are fat tails, that is, there are more large movements in the output gap than is compatible with the normal distribution. That also means that if we were basing our forecasts on the normal distribution, we would underestimate the probability that in any one period a large increase or decrease in the output gap can occur. Finally, the Jarque–Bera test leads to a formal rejection of normality of the movements in the US output gap series.

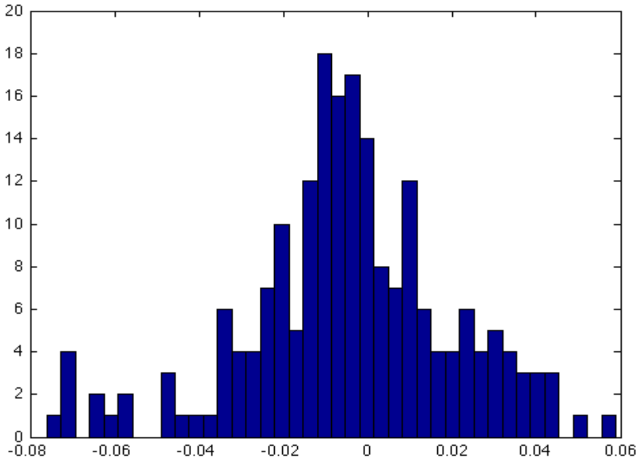

Source: US Department of Commerce and Congressional Budget Office
kurtosis: 3.61; Jarque–Bera: 7.17 with p-value=0.027

Figure 2: Frequency distribution of US output gap (1960–2009)

In this article, I will contrast the rational expectations (Dynamic Stochastic General Equilibrium – DSGE) model with a behavioural macroeconomic model – a model in which agents have cognitive limitations and do not understand the whole picture (the underlying model). I will ask the question of how these two models explain these empirical regularities.

The rational expectations model will be the New Keynesian model. Its characteristic features are price and wage inertia. It is sufficiently well known as not to require much explanation. The behavioural model is less well known, and I will spend more time developing it. Its basic assumption is that agents have cognitive limitations, in other words, they only understand small bits and pieces of the whole model and use simple rules to guide their behaviour. I will introduce rationality in the model through a selection mechanism in which agents evaluate the performance of the rule they are following and decide to switch or to stick to the rule depending on how well the rule performs relative to other rules.

The modelling approach presented here is not the only possible one for modelling agents' behaviour under imperfect information. In fact, a large body of literature has emerged attempting to introduce imperfect information into macroeconomic models. These attempts have been based mainly on the statistical learning approach pioneered by Thomas Sargent (1993) and George Evans and Seppo Honkapohja (2001). This literature leads to important new insights (see, for example, Gaspar, Smets and Vestin, 2006; Orphanides and Williams, 2004; Milani, 2007; and Branch and Evans, 2009). Nevertheless, I feel that this approach still loads individual agents with too many cognitive skills, which they probably do not possess in the real world.[5]

The purpose of this article is to contrast the dynamics of the DSGE model with the behavioural model, and to draw some policy conclusions. It is very much inspired by the new literature on "agent-based macroeconomic models" (see Howitt, 2008; Tesfatsion, 2006; among others). Section 2 presents the behavioural model. The sections that follow it discuss the different implications that the behavioural model has when contrasted with the rational expectations model. Section 7 presents some empirical evidence and concludes with a discussion of some methodological issues.

2. A behavioural macroeconomic model

In this section, the modelling strategy is described by presenting a standard aggregate-demand—aggregate-supply model augmented with a Taylor rule. The novel feature of the model is that agents use simple rules – heuristics – to forecast the future. These rules are subjected to an adaptive learning mechanism, that is, agents endogenously select the forecasting rules that have delivered the highest performance ("fitness") in the past. This selection mechanism acts as a disciplining device on the kind of rules that are acceptable. Since agents use different heuristics, we obtain heterogeneity. This, as will be shown, creates endogenous business cycles.

This behavioural model is contrasted with a similar model that incorporates rational expectations and that is interpreted as a stylized version of DSGE models. This comparison will make it possible to focus on some crucial differences in the transmission of shocks, in particular, of monetary policy shocks.

---

[5] See the fascinating book by Gerd Gigerenzer and Peter Todd (1999) on the use of simple heuristics as compared to statistical (regression) learning.

2.1 The model

The model consists of an aggregate demand equation, an aggregate supply equation and a Taylor rule.

The aggregate demand equation is specified in the standard way, that is,

$$y_t = a_1 \widetilde{E}_t y_{t+1} + (1-a_1) y_{t-1} + a_2 (r_t - \widetilde{E}_t \pi_{t+1}) + \varepsilon_t \tag{1}$$

where y_t is the output gap in period t, r_t the nominal interest rate, π_t the rate of inflation, and ε_t a white noise disturbance term. $\widetilde{E}_t$ is the expectations operator where the tilde above E refers to expectations that are not formed rationally. This process will be specified subsequently. I follow the procedure introduced in DSGE models of adding a lagged output to the demand equation. This is usually justified by invoking habit formation. I keep this assumption here as I want to compare the behavioural model with the DSGE rational expectations model. I will show in Section 4, however, that I do not really need this inertia-building device to generate inertia in the endogenous variables.

The aggregate supply equation can be derived from profit maximization of individual producers. As in DSGE models, a Calvo pricing rule and an indexation rule used in adjusting prices is assumed. This leads to a lagged inflation variable in the equation.[6] The supply curve can also be interpreted as a New Keynesian Philips curve:

$$\pi_t = b_1 \widetilde{E}_t \pi_{t+1} + (1-b_1)\pi_{t-1} + b_2 y_t + \eta_t . \tag{2}$$

Finally, the Taylor rule describes the behaviour of the central bank

$$r_t = c_1(\pi_t - \pi^*) + c_2 y_t + c_3 r_{t-1} + u_t \tag{3}$$

[6] It is now standard in DSGE models to use a pricing equation in which marginal costs enter on the right-hand side. Such an equation is derived from profit maximization in a world of imperfect competition. It can be shown that under certain conditions, the aggregate supply equation (3) is equivalent to such a pricing equation (see Gali, 2008; and Smets and Wouters, 2003).

where π^* is the inflation target, which for the sake of convenience will be set equal to 0. Note that, as is commonly done, the central bank is assumed to smooth the interest rate. This smoothing behaviour is represented by the lagged interest rate in equation (3). Ideally, the Taylor rule should be formulated using a forward-looking inflation variable, that is, central banks set the interest rate on the basis of their forecasts about the rate of inflation. This was not done here in order to maintain simplicity in the model.

Introducing heuristics in forecasting output

Agents are assumed to use simple rules (heuristics) to forecast the future output and inflation. I start with very simple forecasting heuristics and apply it to the forecasting rules of future output. I assume two types of forecasting rules. The first rule can be called "fundamentalist". Agents estimate the steady-state value of the output gap (which is normalized at 0) and use this to forecast the future output gap. (In a later extension, it will be assumed that agents do not know the steady-state output gap with certainty and only have biased estimates of it). The second forecasting rule is an "extrapolative" one. It does not presuppose that agents know the steady-state output gap; they are agnostic about it. Instead, agents extrapolate the previously observed output gap into the future.

The two rules are specified as follows.

The fundamentalist rule is defined by $\widetilde{E}_t^f y_{t+1} = 0$. (4)

The extrapolative rule is defined by $\widetilde{E}_t^e y_{t+1} = y_{t-1}$. (5)

This kind of simple heuristic has often been used in the behavioural finance literature where agents are assumed to use fundamentalist and chartist rules (see Brock and Hommes, 1997; Branch and Evans, 2006; De Grauwe and Grimaldi, 2006). It is probably the simplest possible assumption one can make about how agents, who experience cognitive limitations, use rules that embody limited knowledge to guide their behaviour. In this sense, they are bottom-up rules. They only require agents to use information they understand, and do not require them to understand the whole picture.

Thus, the specification of the heuristics in (4) and (5) should not be interpreted as a realistic representation of how agents forecast. Rather, it is a parsimonious representation of a world where agents do not know the "Truth" (that is, the underlying model). The use of simple rules does not mean that the agents are dumb and that they do not want to learn from their errors. I will specify a learning mechanism later in this section in which these agents continuously try to correct for their errors by switching from one rule to the other.

The market forecast is obtained as a weighted average of these two forecasts, that is,

$$\widetilde{E}_t y_{t+1} = \alpha_{f,t} \widetilde{E}_t^f y_{t+1} + \alpha_{c,t} \widetilde{E}_t^e \tag{6}$$

$$\widetilde{E}_t y_{t+1} = \alpha_{f,t} 0 + \alpha_{c,t} y_{t-1} \tag{7}$$

$$\text{and } \alpha_{f,t} + \alpha_{e,t} = 1 \tag{8}$$

where $\alpha_{f,t}$ and $\alpha_{e,t}$ are the probabilities that agents use a fundamentalist or an extrapolative rule, respectively.

A methodological issue arises here. The forecasting rules (heuristics) introduced are not derived at the micro level and then aggregated. Instead, they are imposed ex post, on the demand and supply equations. This has also been the approach in the learning literature pioneered by Evans and Honkapohja (2001). One could argue, therefore, that my modelling technique is still not fully bottom-up. Ideally one would like to derive the heuristics from the micro level in an environment in which agents experience cognitive problems. Our knowledge about how to model this behaviour at the micro level and how to aggregate it is too sketchy, however, so I have not tried to do so.[7] Clearly, this is an area that will have to be researched in the future.

[7] Psychologists and brains scientists struggle to understand how our brain processes information. There is as yet no generally accepted model we could use to model the micro-foundations of information processing. There have

As indicated earlier, agents are rational in the sense that they continuously evaluate their forecasting performance. I apply notions of discrete choice theory (see Anderson, de Palma, and Thisse, 1992; and Brock and Hommes, 1997) in specifying the procedure agents follow in this evaluation process. Discrete choice theory analyses how agents decide between different alternatives. The theory takes the view that agents are boundedly rational, that is, utility has a deterministic component and a random component. Agents compute the forecast performance of the different heuristics as follows:

$$U_{f,t} = -\sum_{k=0}^{\infty} \omega_k \left[y_{t-k-1} - \widetilde{E}_{f,t-k-2} y_{t-k-1} \right]^2 \tag{9}$$

$$U_{e,t} = -\sum_{k=0}^{\infty} \omega_k \left[y_{t-k-1} - \widetilde{E}_{e,t-k-2} y_{t-k-1} \right]^2 \tag{10}$$

where $U_{f,t}$ and $U_{e,t}$ are the forecast performances (utilities) of the fundamentalists and extrapolators, respectively. These are defined as the mean squared forecasting errors (MSFEs) of the optimistic and pessimistic forecasting rules; ω_k are geometrically declining weights.

Applying discrete choice theory, the probability that an agent will use the fundamentalist forecasting rule is given by the expression (Anderson, de Palma, and Thisse, 1992; Brock and Hommes, 1997):

$$\alpha_{f,t} = \frac{\exp\left(\gamma U_{f,t}\right)}{\exp(\gamma U_{f,t}) + \exp(\gamma U_{e,t})}. \tag{11}$$

Similarly the probability that an agent will use the extrapolative forecasting rule is given by:

$$\alpha_{e,t} = \frac{\exp\left(\gamma U_{e,t}\right)}{\exp(\gamma U_{f,t}) + \exp(\gamma U_{e,t})} = 1 - \alpha_{f,t}. \tag{12}$$

---

been some attempts, however, to provide micro-foundations of models with agents experiencing cognitive limitations See, for example, Kirman (1992) and Delli Gatti, et al. (2005).

Equation (11) says that as the past forecasting performance of the fundamentalists improves relative to that of the extrapolators, agents are more likely to select the fundamentalist rule about the output gap for their future forecasts. As a result, the probability that agents will use the fundamentalist rule increases. Equation (12) has a similar interpretation. The parameter γ measures the "intensity of choice". It parametrizes the extent to which the deterministic component of utility determines actual choice. When $\gamma = 0$, utility is purely stochastic. In that case, agents decide to be fundamentalist or extrapolator by tossing a coin, and the probability to be fundamentalist (or extrapolator) is exactly 0.5. When $\gamma = \infty$, utility is fully deterministic and the probability of using a fundamentalist rule is either 1 or 0. The parameter γ can also be interpreted as expressing a willingness to learn from past performance. When $\gamma = 0$, this willingness is zero; it increases with the size of γ.

Note that this selection mechanism is the disciplining device introduced in this model on the kind of rules of behaviour that are acceptable. Only those rules that pass the fitness test remain in place. The others are weeded out. In contrast with the disciplining device implicit in rational expectations models, implying that agents have superior cognitive capacities, we do not have to make such an assumption here.

It should also be stressed that although individuals use simple rules in forecasting the future, this does not mean that they fail to learn. In fact, the fitness criterion used should be interpreted as a learning mechanism based on "trial and error". When observing that the rule they use performs less well than the alternative rule, agents are willing to switch to the better performing rule. Put differently, agents avoid making systematic mistakes by constantly being willing to learn from past mistakes and to change their behaviour. This also ensures that the market forecasts are unbiased.

The mechanism driving the selection of the rules introduces a self-organizing dynamic into the model. This dynamic goes beyond the capacity of understanding of any one individual in the model. In this sense it is a bottom-up system. It contrasts with the mainstream macroeconomic models in which it is assumed that some or all agents can take a bird's eye view and understand the whole picture. These agents not only understand the whole picture, but also use this whole picture to decide on their optimal behaviour. Thus, there is a one-to-one correspondence between the total information embedded in the world and the individuals' brains.

Introducing heuristics in forecasting inflation

Agents also have to forecast inflation. A similar simple heuristics is used as in the case of output gap forecasting, with one rule that could be called a fundamentalist rule and the other an extrapolative rule (see Brazier et al., 2006, for a similar set-up). The fundamentalist rule is based on the announced inflation target, meaning that agents using this rule have confidence in its credibility and use it to forecast inflation. The extrapolative rule is used by agents who do not trust the announced inflation target. Instead they extrapolate inflation from the past into the future.

The fundamentalist rule will be called an "inflation targeting" rule. It consists in using the central bank's inflation target to forecast future inflation, such that

$$\widetilde{E}_t^{tar} = \pi^* \tag{13}$$

where the inflation target π^* is normalized to be equal to 0.

The "extrapolators" are defined by $\widetilde{E}_t^{ext}\pi_{t+1} = \pi_{t-1}$. (14)

The market forecast is a weighted average of these two forecasts, such that

$$\widetilde{E}_t\pi_{t+1} = \beta_{tar,t}\widetilde{E}_t^{tar}\pi_{t+1} + \beta_{ext,t}\widetilde{E}_t^{ext}\pi_{t+1} \tag{15}$$

or

$$\widetilde{E}_t\pi_{t+1} = \beta_{tar,t}\pi^* + \beta_{ext,t}\pi_{t-1} \tag{16}$$

and $$\beta_{tar,t} + \beta_{ext,t} = 1. \tag{17}$$

The same selection mechanism is used as in the case of output forecasting to determine the probabilities of agents who trust the inflation target and those who do not and revert to extrapolation of past inflation, such that

$$\beta_{tar,t} = \frac{\exp(\gamma U_{tar,t})}{\exp(\gamma U_{tar,t}) + \exp(\gamma U_{ext,t})} \tag{18}$$

$$\beta_{ext,t} = \frac{\exp\left(\gamma U_{ext,t}\right)}{\exp(\gamma U_{tar,t}) + \exp(\gamma U_{ext,t})} \tag{19}$$

where $U_{tar,t}$ and $U_{ext,t}$ are the weighted averages of past squared forecast errors of using targeter and extrapolator rules, respectively. These are defined in the same way as in (9) and (10).

This inflation forecasting heuristics can be interpreted as a procedure of agents to find out how credible the central bank's inflation targeting is. If this is very credible, using the announced inflation target will produce good forecasts and as a result, the probability that agents will rely on the inflation target will be high. If, on the other hand, the inflation target does not produce good forecasts (compared to a simple extrapolation rule), the probability that agents will use it will be small.

The solution of the model is found by first substituting (3) into (1) and rewriting in matrix notation. This yields:

$$\begin{bmatrix} 1 & -b_2 \\ -a_2c_1 & 1-a_2c_2 \end{bmatrix} \begin{bmatrix} \pi_t \\ y_t \end{bmatrix} = \begin{bmatrix} b_1 & 0 \\ -a_2 & a_1 \end{bmatrix} \begin{bmatrix} \tilde{E}_t\pi_{t+1} \\ \tilde{E}_t y_{t+1} \end{bmatrix} + \begin{bmatrix} 1-b_1 & 0 \\ 0 & 1-a_1 \end{bmatrix} \begin{bmatrix} \pi_{t-1} \\ y_{t-1} \end{bmatrix} + \begin{bmatrix} 0 \\ a_2c_3 \end{bmatrix} r_{t-1} + \begin{bmatrix} \eta_t \\ a_2u_t + \varepsilon_t \end{bmatrix}$$

or

$$\boldsymbol{A}\ \boldsymbol{Z_t} = \boldsymbol{B}\ \boldsymbol{\tilde{E}_t Z_{t+1}} + \boldsymbol{C}\ \boldsymbol{Z_{t-1}} + \boldsymbol{b}\ r_{t-1} + \boldsymbol{v_t} \tag{20}$$

where bold characters refer to matrices and vectors. The solution for $\boldsymbol{Z_t}$ is given by

$$\boldsymbol{Z_t} = \boldsymbol{A^{-1}} \left[\boldsymbol{B}\ \boldsymbol{\tilde{E}_t Z_{t+1}} + \boldsymbol{C}\ \boldsymbol{Z_{t-1}} + \boldsymbol{b}\ r_{t-1} + \boldsymbol{v_t}\right]. \tag{21}$$

The solution exists if the matrix $\boldsymbol{A}$ is non-singular, that is, if $(1-{}_2c_2)a_2b_2c_1 \neq 0$. The system (21) describes the solution for y_t and π_t given the forecasts of y_t and π_t. The latter have been specified in equations (4) to (12) and can be substituted into (21). Finally, the solution for r_t is found by substituting y_t and π_t obtained from (21) into (3).

My research strategy consists in comparing the dynamics of this behavioural model with the same structural model (aggregate demand equation (1), aggregate supply equation (2) and Taylor rule equation (3)) under rational expectations, which I interpret as a stylized DSGE model.

The model consisting of equations (1) to (3) can be written in matrix notation as follows:

$$\begin{bmatrix} 1 & -b_2 & 0 \\ 0 & 1 & -a_2 \\ -c_1 & -c_2 & 1 \end{bmatrix}\begin{bmatrix} \pi_t \\ y_t \\ r_t \end{bmatrix} = \begin{bmatrix} b_1 & 0 & 0 \\ -a_2 & a_1 & 0 \\ 0 & 0 & 0 \end{bmatrix}\begin{bmatrix} E_t\pi_{t+1} \\ E_t y_{t+1} \\ E_t r_{t+1} \end{bmatrix} + \begin{bmatrix} 1-b_1 & 0 & 0 \\ 0 & 1-a_1 & 0 \\ 0 & 0 & a_3 \end{bmatrix}\begin{bmatrix} \pi_{t-1} \\ y_{t-1} \\ r_{t-1} \end{bmatrix} + \begin{bmatrix} \eta_t \\ \varepsilon_t \\ u_t \end{bmatrix}$$

$$\Omega\, \boldsymbol{Z_t} = \Phi\, \boldsymbol{E_t Z_{t+1}} + \Lambda\, \boldsymbol{Z_{t-1}} + \boldsymbol{v_t} \tag{22}$$

$$\boldsymbol{Z_t} = \Omega^{-1}\left[\Phi\, \boldsymbol{E_t Z_{t+1}} + \Lambda\, \boldsymbol{Z_{t-1}} + \boldsymbol{v_t}\right] \tag{23}$$

This model can be solved under rational expectations using the Binder–Pesaran (1996) procedure.

2.2 Calibrating the model

I proceed by calibrating the model. In the Appendix, the parameters used in the calibration exercise are presented. The model was calibrated in such a way that the time units can be considered to be months. A sensitivity analysis of the main results to changes in some of the parameters of the model will be presented. The three shocks (demand, supply and interest rate) are independent and identically distributed (i.i.d.) with standard deviations of 0.5 per cent.

3. Animal spirits, learning and forgetfulness

In this section, simulations of the behavioural model in the time domain are presented and interpreted. The upper panel of Figure 3 below shows the time pattern of the output gap produced by the behavioural model. A strong cyclical movement in the output gap can be observed. The lower panel of Figure 3 shows a variable called "animal spirits".[8] It represents the evolution of the fractions of the agents who

[8] See Nuti (2009) on the different interpretations of "Animal Spirits". The *locus classicus* is Keynes (1936). See also Farmer (2006) and the recent book by George Akerlof and Robert Shiller (2009).

extrapolate a positive output gap. Thus, when the curve reaches +1, all agents are extrapolating a positive output gap; when the curve reaches 0, no agents are extrapolating a positive output gap. In fact, in that case they all extrapolate a negative output gap. The curve thus shows the degree of optimism and pessimism of agents who make forecasts of the output gap.

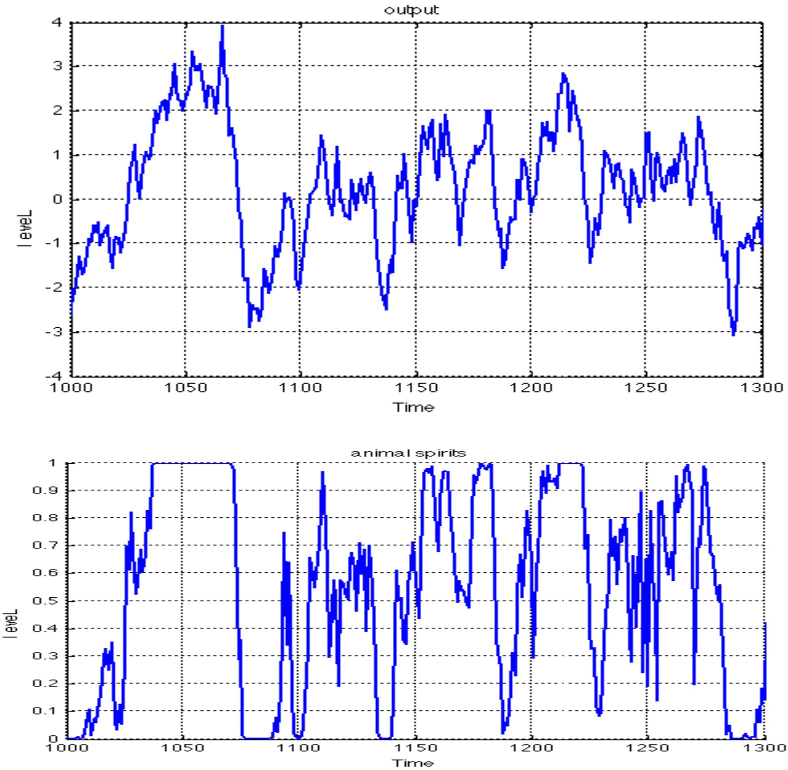

Figure 3: Output gap in behavioural model

Combining the information of the two panels in Figure 3, it can be seen that the model generates endogenous waves of optimism and pessimism. During some periods optimists (that is, agents who extrapolate positive output gaps) dominate, and this translates into above-average output growth. These optimistic periods are followed by pessimistic ones when pessimists (that is, agents who extrapolate negative output gaps) dominate, and the growth rate of output is below average.

These waves of optimism and pessimism are essentially unpredictable. Other realizations of the shocks produce different cycles with the same general characteristics.

These endogenously generated cycles in output are made possible by a self-fulfilling mechanism that can be described as follows. A series of random shocks creates the possibility that one of the two forecasting rules – say the extrapolating one – delivers a higher payoff, that is, a lower mean squared forecast error (MSFE). This attracts agents that were using the fundamentalist rule. If the successful extrapolation happens to be a positive extrapolation, more agents will start extrapolating the positive output gap. The "contagion effect" leads to an increasing use of the optimistic extrapolation of the output gap, which in turn stimulates aggregate demand. Optimism is therefore self-fulfilling. A boom is created. At some point, negative stochastic shocks and/or the reaction of the central bank through the Taylor rule make a dent in the MSFE of the optimistic forecasts. Fundamentalist forecasts may become attractive again, but it is equally possible that pessimistic extrapolation becomes attractive and therefore fashionable again. The economy turns around.

These waves of optimism and pessimism can be understood to be searching (learning) mechanisms of agents who do not fully understand the underlying model but are continuously searching for the truth. An essential characteristic of this searching mechanism is that it leads to systematic correlation in beliefs (for example, optimistic or pessimistic extrapolations). This systematic correlation is at the core of the booms and busts created in the model. Note, however, that when computed over a significantly long period of time, the average error in the forecasting goes to zero. In this sense, the forecast bias tends to disappear asymptotically.

The results concerning the time path of inflation are shown in Figure 4. The lower panel shows the fraction of agents using the extrapolator heuristics, that is, the agents who do not trust the inflation target of the central bank. One can identify two regimes. There is a regime in which the fraction of extrapolators fluctuates around 50 per cent, which also implies that the fraction of forecasters using the inflation target as their guide (the "inflation targeters") is around 50 per cent. This is sufficient to maintain the rate of inflation within a narrow band of approximately + or – 1 per cent around the central bank's inflation target. There is a second regime though that occurs when the extrapolators are dominant. During this regime, the rate of inflation fluctuates significantly more. Thus, the inflation targeting of the central bank is fragile. It can be undermined when forecasters decide that relying on past inflation movements produces better forecasting performances than relying on the central bank's inflation target. This can occur unpredictably as a result of stochastic shocks in supply and/or demand. I will return to the question of how the central bank can reduce this loss of credibility.

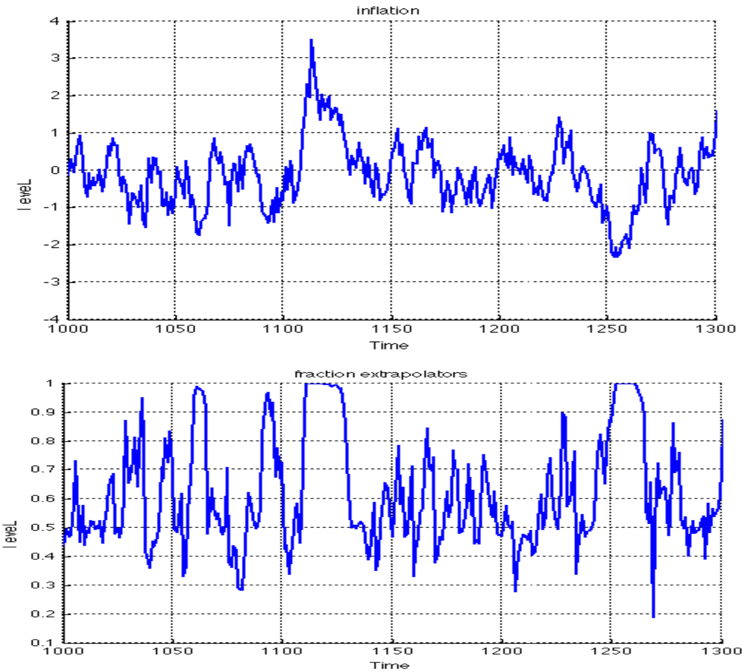

Figure 4: Inflation in behavioural model

The simulations reported in the previous section assumed a given set of numerical values of the parameters of the model. It was found that for this set of parameter values, animal spirits (measured by the movements in the fraction of optimistic extrapolators) emerge and affect the fluctuations of the output gap. The correlation coefficient between the fraction of optimists and the output gap in the simulation reported in Figure 3 is 0.86. One would like to know how this correlation evolves when one changes the parameter values of the model. I concentrate on two parameter values here, the intensity of choice parameter, γ, and the memory agents have when calculating the performance of their forecasting. The latter is represented by the parameter ω_k in equations (9) and (10) and is a series of declining weights attached to past forecasting errors. I define $\omega_k = (1-\rho)\rho^k$ (and $0 \leq \rho \leq 1$). The parameter ρ can then be interpreted as a measure of the memory of agents. When $\rho = 0$, there is no memory, meaning only last period's performance matters in evaluating a forecasting rule; when $\rho = 1$, there is infinite memory, meaning all past errors, however far in the past, obtain the same weight.

The results of the sensitivity analysis are shown in Figure 5. The left-hand panel shows the correlation between the output gap and the fraction of optimistic extrapolators (animal spirits) for increasing values of the intensity of choice parameter, γ. It can be seen that when γ is zero (that is, when the switching mechanism is purely stochastic), this correlation is zero. The interpretation is that in an environment in which agents decide purely randomly – in other words, they do not react to the performance of their forecasting rule – there are no systematic waves of optimism and pessimism (animal spirits) that can influence the business cycle. When γ increases, the correlation increases sharply. Thus, in an environment in which agents learn from their mistakes, animal spirits arise. One thus needs a minimum level of rationality (in the sense of a willingness to learn) for animal spirits to emerge and to influence the business cycle. Figure 3 shows that this is achieved with relatively low levels of γ.

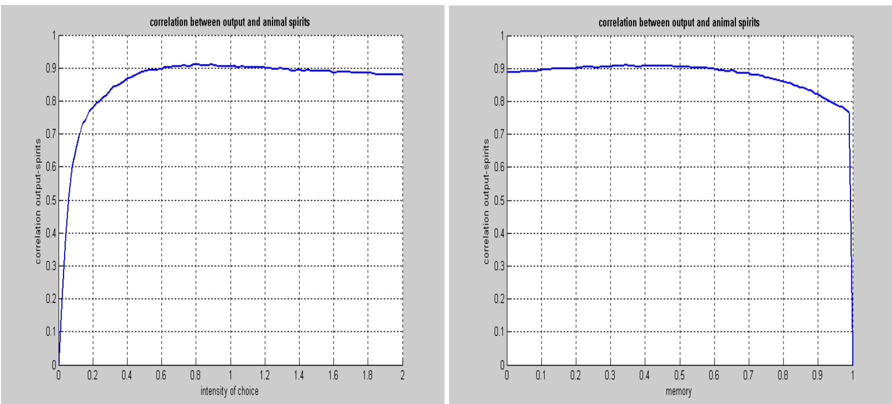

Figure 5: Correlations between output gap and fraction of optimists

The right-hand panel shows the correlation between the output gap and the fraction of optimists for increasing values of the memory parameter ρ. It can be seen that when $\rho = 1$ the correlation is zero. This is the case where agents attach the same weight to all past observations, however far in the past they occurred. Put differently, when agents have infinite memory, they forget nothing. In that case, animal spirits do not occur. Thus one needs some forgetfulness (which is a cognitive limitation) to produce animal spirits. Note that the degree of forgetfulness does not have to be large. For values of ρ below 0.98, the correlations between output and animal spirits are quite high.

Having presented the main features of the behavioural model, I will now proceed to show how this model leads to a view of macroeconomic dynamics that contrasts greatly with the one obtained from the rational-expectations DSGE models. I will concentrate on two areas. The first one has to do with the business cycle theories implicit in the behavioural and the rational expectations models. The second one focuses on the implications for monetary policies.

4. Two different business cycle theories

Are the behavioural and the New-Keynesian models capable of mimicking these empirical regularities? Let us first focus on the behavioural model presented in the previous section. Figure 3 presented a typical simulation of the output gap obtained in that model. The autocorrelation coefficient of the output gap obtained from Figure 3 is 0.95, which is very close to 0.94, that is, the autocorrelation of the output gap in the USA during 1960–2009 (see the introduction). In addition, my behavioural macroeconomic model produces movements of output that are very different from the normal distribution. I show this by presenting the histogram of the output gaps obtained in Figure 3. The result is presented in Figure 6. The frequency distribution of the output gap deviates significantly from a normal distribution. There is excess kurtosis (kurtosis $=$ 4.4), meaning there is too much concentration of observations around the mean for the distribution to be normal. In addition, there are fat tails. This means that there are too many observations that are extremely small or extremely large to be compatible with a normal distribution. I also applied a more formal test of normality, the Jarque–Bera test, which rejected normality. Note that the non-normality of the distribution of the output gap is produced endogenously by the model, as I feed the model with normally distributed shocks.

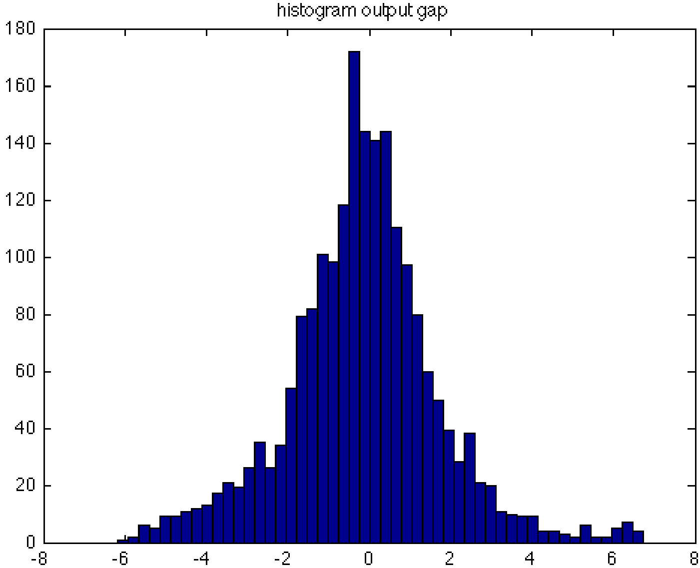

Kurtosis=4.4, Jarque-Bera = 178.4 (p-value = 0.001)

Figure 6: Frequency distribution of simulated output gap

This result is not without implications. It implies that when we use the assumption of normality in macroeconomic models, we underestimate the probability of large changes. In this particular case, assuming normal distributions tends to underestimate the probability that intense recessions or booms will occur. The same is true in finance models that assume normality. These models greatly underestimate the probability of extremely large asset price changes. In other words, they underestimate the probability of large bubbles and crashes. To use the metaphor introduced by Nassim Taleb, there are many more Black Swans than theoretical models based on the normality assumption.

It is fine to observe this phenomenon. It is even better to have an explanation for it. My model provides such an explanation. It is based on the particular dynamics of "animal spirits", illustrated in Figure 7, which shows the frequency distribution of the animal spirits index (defined earlier). This index is associated with the frequency distribution of the output gap obtained in Figure 6. From Figure 7, we observe that there is a concentration of the animal spirits at the extreme values of 0 and 1 and also in the middle of the distribution (but more spread out). This feature provides the key explanation of the non-normality of the movements of the output gap.

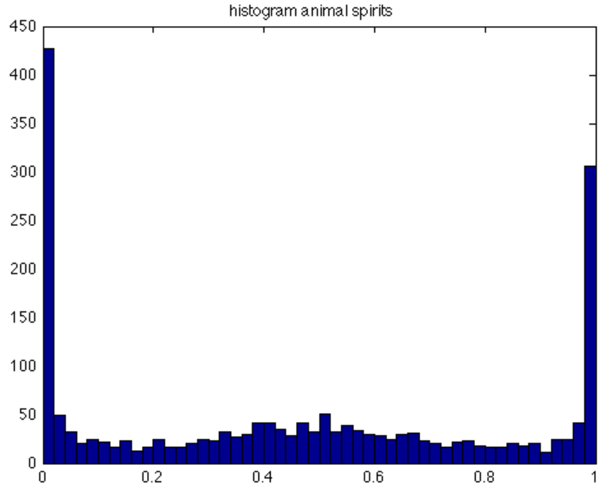

Figure 7: Frequency distribution simulated animal spirits

When the animal spirits index clusters in the middle of the distribution, we have tranquil periods. There is no particular optimism or pessimism, and agents use a fundamentalist rule to forecast the output gap. At irregular intervals, however, the economy is gripped by either a wave of optimism or of pessimism. The nature of these waves is that beliefs get correlated. Optimism breeds optimism; pessimism breeds pessimism. This can lead to situations where everybody has become either optimistic or pessimistic. These periods are characterized by extreme positive or negative movements in the output gap (booms and busts).

From the previous discussion, it follows that my behavioural macroeconomic model has a strong prediction about how the movements of the output gap are distributed. These movements should be non-normal. This is also what one observes in reality.

How well does the New Keynesian (DSGE) model perform in mimicking the empirical regularities about the business cycle? I simulated the Rational Expectations version of equations (1) to (3) (the New Keynesian model) using the same calibration. I show the movements of the simulated output gap in Figure 8. The upper panel shows the output gap in the time domain and the lower panel in the frequency domain. The autocorrelation in the output gap is 0.77, which is significantly lower than in the observed data (for the USA, I found 0.94). In addition, these output gap movements are normally distributed (see lower panel). We cannot reject that the distribution is normal.

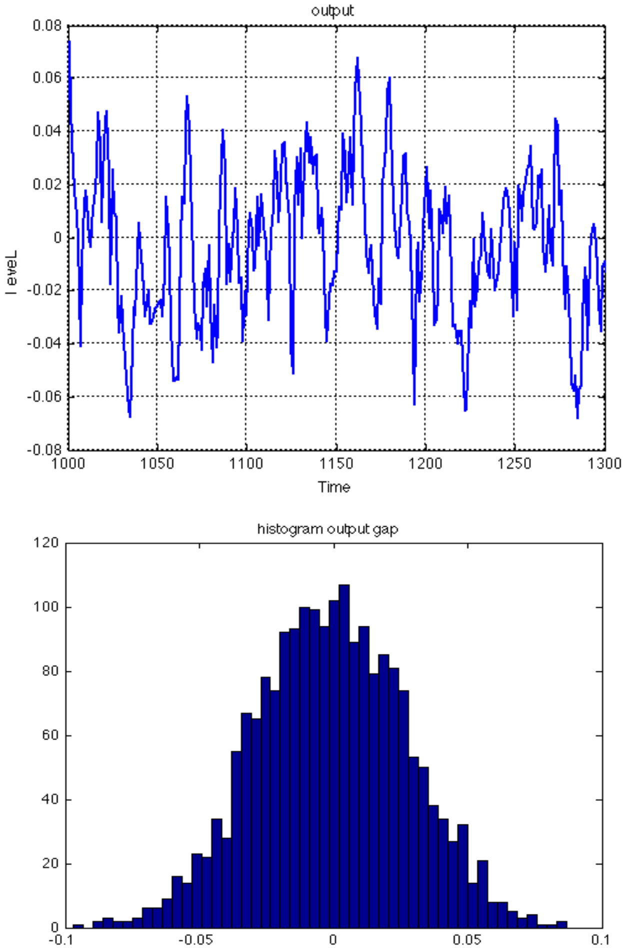

kurtosis: 2.9; Jarque-Bera: 1.03 with p-value=0.5

Figure 8: Simulated output gap in extended New Keynesian model

The next step in making this model more empirically relevant consists in adding autocorrelation in the error terms. This is now the standard procedure in DSGE models (see Smets and Wouters, 2003). I have done the same with my version of the New Keynesian model and assumed that the autocorrelation of the error terms in equations (1) to (3) is equal to 0.9. The result of this assumption is shown in the simulations of the output gap in Figure 9 below. We now obtain movements of the output gap that resemble real-life movements. The autocorrelation of the output gap is now 0.98, which is very close to the observed number of 0.94 in the postwar US output gap. We still cannot reject normality though (see the Jarque–Bera test). This is a problem that DSGE models have not been able to solve.

Thus, in order to mimic business cycle movements, the New Keynesian (DSGE) model builders have had recourse to introducing autocorrelation in the error terms (the shocks that hit the economy). This trick has allowed DSGE models to closely fit observed data (see Smets and Wouters, 2003). This success has been limited to the first and second moments of the movements of output, but not to the highest moments (kurtosis, fat tails). The latter failure has the implication that in order to explain a large movement in output (for example, a deep recession, or a strong boom), DSGE models have to rely on large unpredictable shocks.

There are two problems with this theory of the business cycle implicit in the DSGE models. First, business cycles are not the result of endogenous dynamics. They occur as a result of exogenous shocks and slow transmission of those shocks. Put differently, the DSGE models picture a world populated by rational agents who are fully informed. In such a world, there would never be business cycles. The latter arise because of exogenous disturbances and of constraints on agents' ability to react instantaneously to these shocks. A given shock will thus produce ripple effects in the economy, that is, cyclical movements. The second problem is methodological. When the New Keynesian model is tested empirically, the researcher finds that a lot of the output dynamics are not predicted by the model. These unexplained dynamics are found in the error term. Everything is fine up to this point. The next step taken by DSGE modellers is to conclude that these errors (typically autocorrelated) should be considered as exogenous shocks.

The problem with this approach is that it is not scientific. When the DSGE modeller finds dynamics that are not predicted by the model, he or she decides that the New Keynesian model must nevertheless be right (because there can be no doubt that individual agents are rational), and thus that the deviation between the observed dynamics and those predicted by the model must come from outside the model.

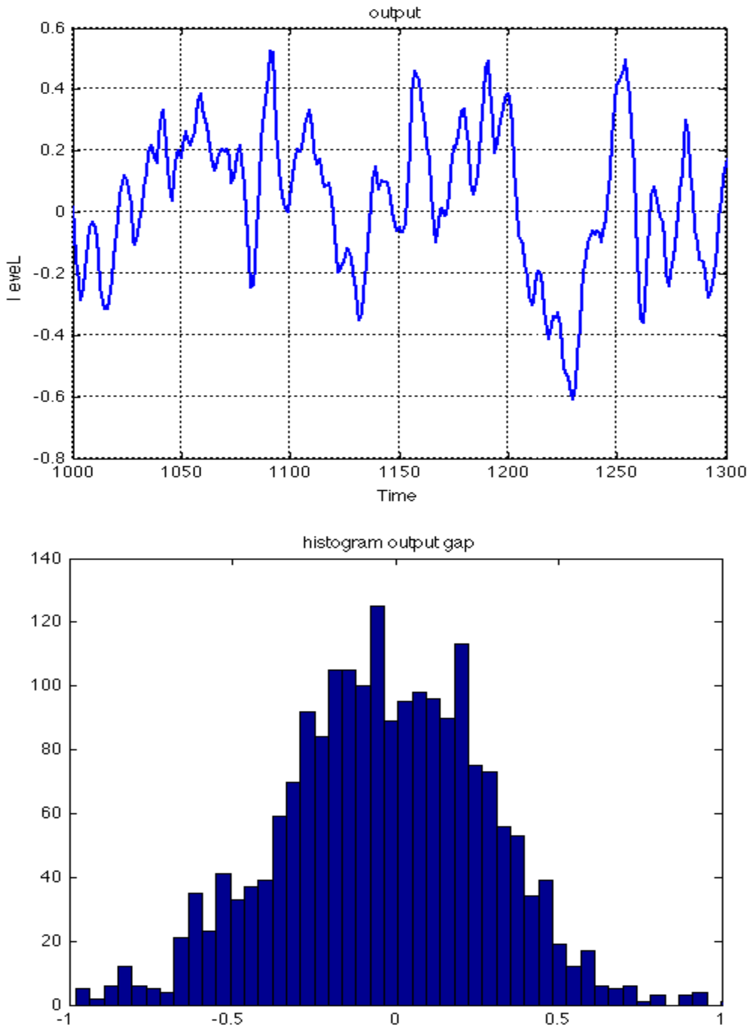

kurtosis: 3.16; Jarque-Bera: 3.2 with p-value=0.17

Figure 9 : Simulated output gap in extended New Keynesian model and autocorrelated errors

5. The role of output stabilization

Modern macroeconomics, in general, and DSGE models, in particular, have provided the intellectual foundation of inflation targeting. Until the eruption of the financial crisis in 2007, inflation targeting strategies had become the undisputed policy framework modern central banks should adopt. And most did. The official holders of macroeconomic wisdom declared that this step towards inflation targeting constituted a great victory of macroeconomics as a science (Woodford, 2009). From now on, we would be living in a more stable macroeconomic environment – a "Great Moderation". How things can change so quickly.

Inflation targeting, of course, does not imply that there is no role for output stabilization. DSGE modellers who have put a New Keynesian flavour into their models have always stressed that wage and price rigidities provide a rationale for output stabilization by central banks (see Clarida et al., 1999; and Gali, 2008). This idea has found its reflection in "flexible" inflation targeting (Svensson, 1997; Woodford, 2003). Because of the existence of rigidities, a central bank should not attempt to keep inflation close to its target all the time. When sufficiently large shocks occur that lead to departures of inflation from its target, the central bank should follow a strategy of gradual return of inflation to its target. The rationale is that in a world of wage and price rigidities, overly abrupt attempts to bring back inflation to its target would require such high increases in the interest rate as to produce overly strong declines in output.

Output stabilization in the DSGE world, however, is very much circumscribed. The need to stabilize arises because of the existence of rigidities in prices that makes it necessary to spread out price movements over longer periods. The limited scope for output stabilization is based on a model characterized by a stable equilibrium. There is no consideration of the possibility that the equilibrium may be unstable or that fluctuations in output have a different origin than price rigidities. Should the scope for output stabilization be enlarged? In order to shed some light on this issue, I will now derive the tradeoff between output and inflation variability in the context of the behavioural model, and formulate some policy conclusions.

The tradeoffs are constructed as follows. The model was simulated 10,000 times, and the average output and inflation variabilities were computed for different values of the Taylor rule parameters. Figure 10 shows how output variability (panel a) and inflation variability (panel b) change as the output coefficient (c_2) in the

Taylor rule increases from 0 to 1. Each line represents the outcome for different values of the inflation coefficient (c_1) in the Taylor rule.

Panel a

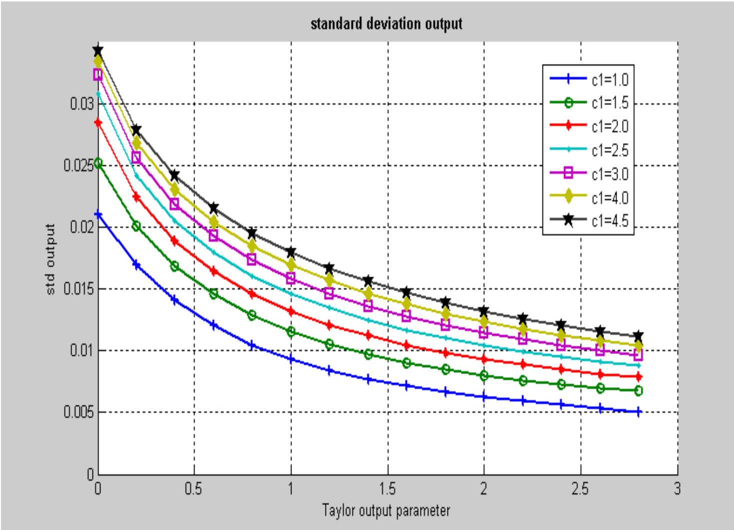

Panel b

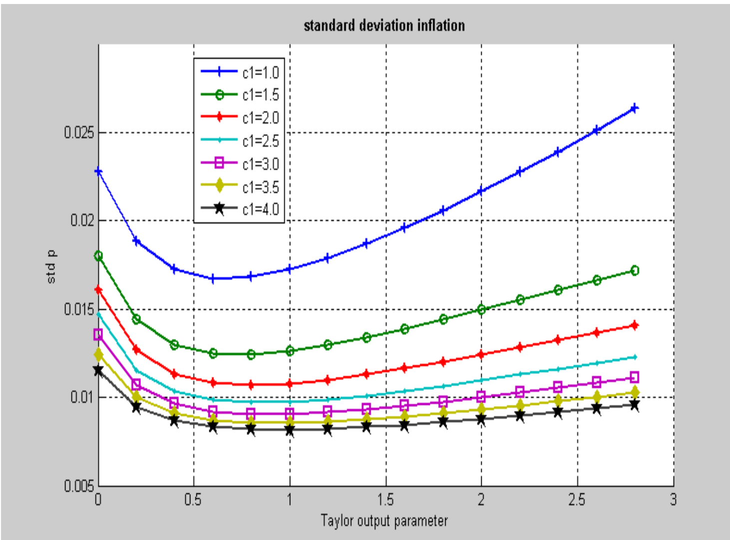

Figure 10: Output and inflation variability

"Panel a" showing the evolution of output variability exhibits the expected result: as the output coefficient (c_2) increases (inflation targeting becomes less strict), output variability tends to decrease. One would now expect that this decline in output variability resulting from more active stabilization comes at the cost of more inflation variability. This, however, is not found in "Panel b". We observe that the relationship is non-linear. As the output coefficient is increased from zero, inflation variability first declines. Only when the output coefficient increases beyond a certain value (in a range from 0.6–0.8) does inflation variability start to increase. Thus, the central bank can reduce both output and inflation variability when it moves away from strict inflation targeting ($c_2 = 0$) and engages in some output stabilization. Not too much though. Too much output stabilization reverses the relationship and increases inflation variability.

Figure 10 makes it possible to construct the tradeoffs between output and inflation variability. These are shown in Figure 11 for different values of the inflation parameter c_1. Take the tradeoff AB. This is the one obtained for $c_1 = 1$. Start from point A on the tradeoff. In point A, the output parameter is $c_2 = 0$ (strict inflation targeting). As output stabilization increases, it first moves downwards. Thus, increased output stabilization by the central bank reduces output and inflation variability. The relation is non-linear, however. At some point, with an overly high output stabilization parameter, the tradeoff curve starts to increase, becoming a "normal" tradeoff: a lower output variability is obtained at the cost of increased inflation variability.

How can we interpret these results? Let us start from the case of strict inflation targeting, that is, the authorities set $c_2=0$. There is no attempt at stabilizing output at all. The ensuing output variability intensifies the waves of optimism and pessimism (animal spirits), which in turn feed back into output volatility. These large waves lead to higher inflation variability. Thus, some output stabilization is good; it reduces both output and inflation variability by preventing overly large swings in animal spirits. With no output stabilization at all ($c_2=0$), the forces of animal spirits are so high that the high output variability also increases inflation volatility through the effect of the output gap on inflation (supply equation). Too much output stabilization, however, reduces the stabilization bonus provided by a credible inflation target. When the central bank attaches too much importance to output stabilization, it creates more scope for better forecasting performance of the inflation extrapolators, leading to more inflation variability.

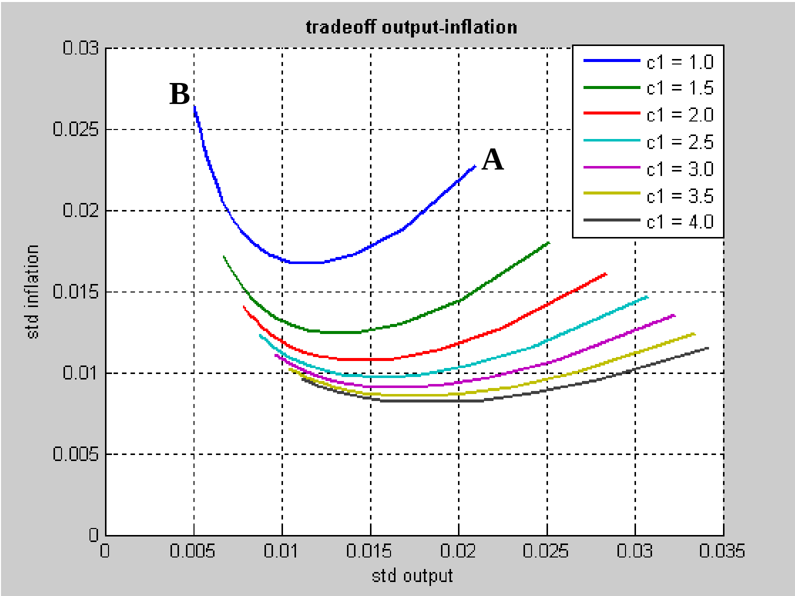

Figure 11: Tradeoffs in the behavioural model

Figure 11 also tells us something important about inflation targeting. We note that increasing the inflation parameter in the Taylor rule (c_1) has the effect of shifting the tradeoffs downwards, in other words, the central bank can improve the tradeoffs by reacting more strongly to changes in inflation.[9] The central bank achieves this improvement in the tradeoff because by reacting more intensely to changes in inflation, it reduces the probability that inflation extrapolators will tend to dominate the market. As a result, it reduces the probability that inflation targeting will lose credibility. Such a loss of credibility destabilizes both inflation and output. Thus, maintaining credibility of inflation targeting is an important source of macroeconomic stability in my behavioural model.

[9] A similar result on the importance of strict inflation is also found in Gaspar, Smets and Vestin (2006), which uses a macromodel with statistical learning.

6. Fiscal policy multipliers: How much do we know?

Since the eruption of the financial crisis in 2007–08, governments of major countries have applied massive policies of fiscal stimulus. This has led to a heated debate about the size of the fiscal policy multipliers. This debate has revealed (once more) how divergent economists' views are about the size of these multipliers (see Wieland, 2010). The estimates of the short-term multipliers vary from 0 to numbers far exceeding 1. There has been a lot of soul-searching about the reasons for these widely divergent estimates.

An important source of these differences is to be found in the use of different models that embody different priors. For example, in mainstream macroeconomic models that incorporate agents with rational expectations (both New Classical and the New Keynesian), fiscal policy multipliers are likely to be very small, as these models typically have Ricardian equivalence embedded in them. That means that agents who anticipate future tax increases following a fiscal stimulus (budget deficit) will start saving more (consuming less) so that one dollar of government spending is offset by 1 dollar of less private spending. In these models, the fiscal policy multiplier is close to zero. In Keynesian models, there is scope for a net stimulatory effect of fiscal policies. Thus, the different estimates of fiscal policy multipliers are not "neutral estimates", but reflect theoretical priors and beliefs that have been put in these models in the construction stage.

My behavioural model makes it possible to shed some additional light on the uncertainty surrounding the effects of fiscal policies. I will do this by studying how a positive shock in aggregate demand produced by a fiscal expansion affects output. I will not give an exhaustive analysis of fiscal policies. The model does not give sufficient detail of government spending and taxation to be able to do that. I will model a fiscal policy shock just as a shock in the demand equation. The model then allows me to establish the nature of uncertainty surrounding such a shock, even in an extremely simple model.

I assume the fiscal policy expansion to occur under two different monetary policy regimes. In the first regime, I assume that the central bank uses the standard Taylor rule as specified in equation (3). Thus, under this regime, the fiscal policy expansion will automatically lead the central bank to raise the interest rate. This follows from the fact that the demand stimulus produces an increase in output and inflation to which the central bank reacts by raising the interest rate.

In the second regime, I assume that the central bank does not react to the stimulus-induced expansion of output and inflation by raising the interest rate. I do this, not because it is realistic, but rather to estimate the pure Keynesian multiplier effect of a fiscal stimulus. The Keynesian multiplier is usually estimated under the assumption of a constant interest rate so that crowding out does not occur.

The results of this fiscal policy stimulus under the two monetary policy regimes are presented in Figure 12 below. The upper two panels show the impulse responses under the two monetary policy regimes. The instantaneous effects of the fiscal stimulus are the same under the two regimes. Under the variable interest rate regime, however, the positive effects of the fiscal stimulus decline faster and undershoot in the negative region more than under the constant interest regime. This is not surprising as under the variable interest rate regime we see that the interest rate is raised substantially (see bottom panel), leading to a quick crowding out.

A second important difference concerns the degree of uncertainty about the size of the output effects of a fiscal stimulus. As the upper panels show, the divergence in the impulse responses is larger in the constant interest rate regime than in the variable interest rate regime. This is also illustrated in the second panels. These show the frequency distribution of the short-term output responses under the two regimes. We observe a wider spread of these short-term output responses under the fixed interest rate regime. The reason is to be found in the fact that animal spirits behave differently under the two monetary regimes. The interest rate response under the variable interest rate regime tends to reduce the impact of animal spirits on the transmission mechanism, thereby reducing the volatility in this transmission. Put differently, when, as a result of the fiscal expansion, the central bank raises the interest rate, it lowers the expansionary effect of this expansion, making it less likely that positive animal spirits will enhance the fiscal policy stimulus.

These results make clear that there is likely to be a great amount of uncertainty about the size of the output effects of fiscal policies. This uncertainty is even more pronounced in the Keynesian scenario of a constant interest rate. This is also the scenario usually associated with the occurrence of a liquidity trap (a horizontal LM-curve). This is the assumption that tends to make fiscal policies most effective. In my model, it is also the assumption that makes the uncertainty about the size of these effects the greatest.

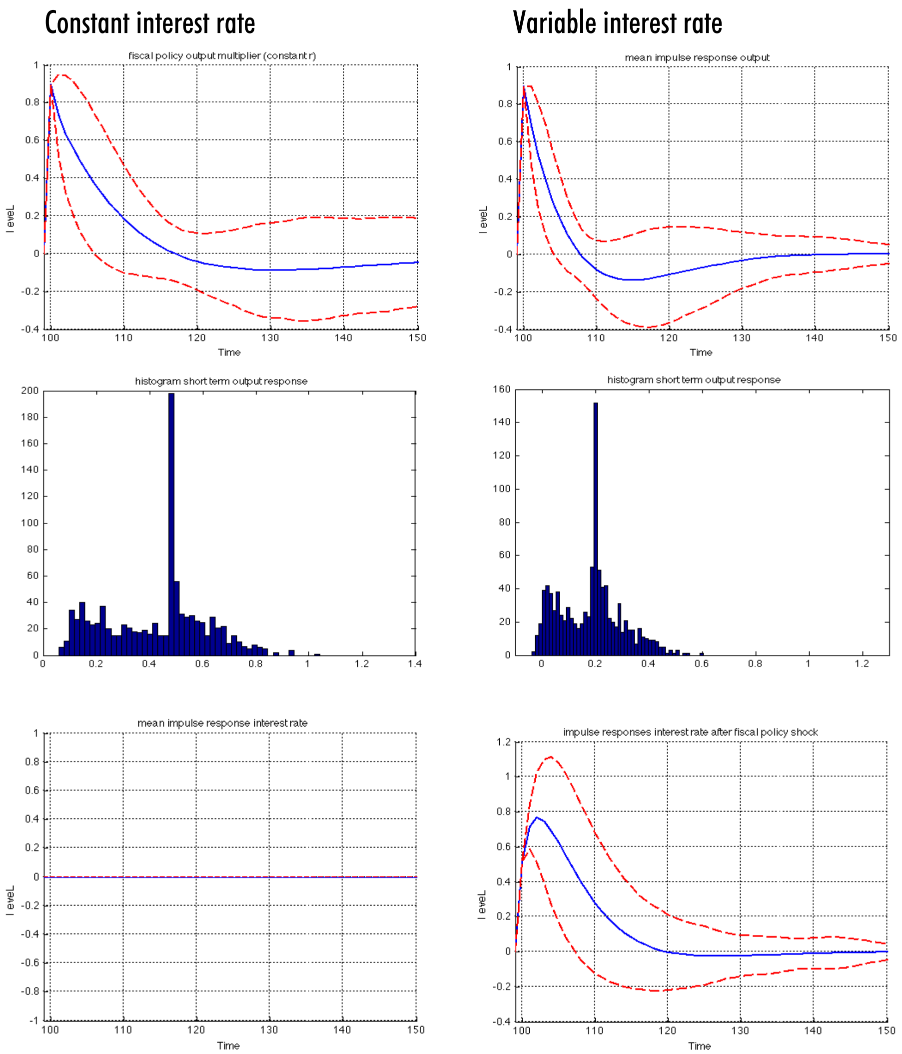

Figure 12: Constant interest rate vs. variable interest rate

These differences are also made clear from a comparison of the long-term fiscal policy multipliers obtained from the same simulations as in Figure 12. The fiscal policy shock underlying the previous simulations is a one-period increase in demand (by one standard deviation). (The closest example of such a shock is the "Cash for Clunkers" car-buying stimulus programmes introduced in many European countries and in the USA in 2009). This temporary increase then produces the impulse responses as given in Figure 12. In order to obtain the long-term multipliers, I add up all the output increases (and declines) following this temporary fiscal policy shock. These long-term fiscal policy multipliers are presented in Figure 13 under the two monetary policy regimes.

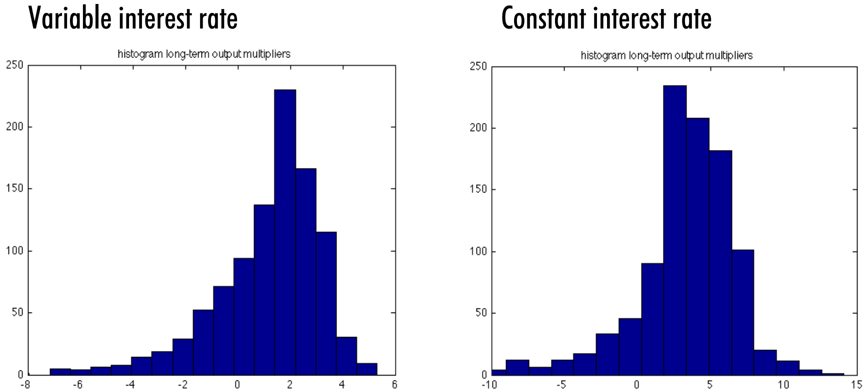

Figure 13: Long-term fiscal policy multipliers: Frequency distribution

Two results stand out. First, as expected, the long-term fiscal policy multipliers are higher under the constant interest rate rule than under the variable interest rate rule. Second, the uncertainty surrounding these long-term multipliers is considerable. And this uncertainty is the most pronounced under the constant interest-rate rule.

It should be stressed again that the nature of the uncertainty here is not the uncertainty surrounding the parameters of the model. I assume exactly the same parameters in all these simulations. Put differently, it is not the uncertainty produced by the use of different models with different prior beliefs about the effectiveness of fiscal policies that yields uncertainty. The uncertainty is due to differences in initial conditions (market sentiments). These differences in market sentiments have a pronounced effect on how the same fiscal policy shock is transmitted in the economy.

7. Conclusion

Capitalism is characterized by booms and busts, in other words, economic activity is often subjected to strong growth followed by sharp declines. As a result, the frequency distribution of the output gap (and output growth) is non-normal, exhibiting excess kurtosis and fat tails. The latter means that if we are basing our forecasts on the normal distribution, we will tend to underestimate the probability that in any one period a large increase or decrease in the output gap can occur.

In this article, I used two alternative models to explain this empirical regularity. One model is the DSGE model, which assumes rational expectations. The other is a behavioural model. The latter is a model in which agents experience cognitive limitations. These limitations force agents to use simple rules to forecast output and inflation. Rationality is introduced into this model by assuming a learning mechanism that allows for the selection of those rules that are more profitable than others.

In the DSGE model, large booms and busts can only be explained by large exogenous shocks. Price and wage rigidities then lead to wavelike movements of output and inflation. Thus, booms and busts are explained exogenously. The fat tails observed in the frequency distribution of the output gap arise because there are large shocks hitting the economy.

My behavioural model provides a very different explanation. The behavioural model creates correlations in beliefs, which in turn generate waves of optimism and pessimism. Such waves produce endogenous cycles, which are akin to the Keynesian animal spirits. Occasionally this correlation of beliefs leads to extreme optimism (explaining booms) followed by extreme pessimism (explaining busts). The behavioural model thus provides for an endogenous explanation of business cycle movements.

In both models, the inflation targeting regime turns out to be of great importance for stabilizing the economy. In the behavioural model, this follows from the fact that credible inflation targeting also helps to reduce correlations in beliefs and the ensuing self-fulfilling waves of optimism and pessimism. Nevertheless, and this is where the behavioural model departs from the rational expectations model, strict inflation targeting is not an optimal policy. Some output stabilization (given a credible inflation target) also helps to reduce the correlation of biased beliefs, thereby reducing the scope for waves of optimism and pessimism to emerge and to destabilize output and inflation.

The behavioural model proposed in this article can be criticized for being "ad hoc". There is no doubt that the model has ad-hoc features, that is, assumptions that cannot be grounded on some deeper principle, and therefore have to be taken for granted. In defence of this "ad-hocness", the following should be stressed. Once we leave the comfortable world of agents who experience no limits to their cognitive abilities, ad-hoc assumptions are inevitable. This is due to the fact that we do not fully comprehend the way individuals with cognitive limitations process information. In contrast, there is no secret in how the superbly informed individuals in the rational expectations world process information. They understand the model, and therefore there is only one way to write down how they form their expectations. This feature may give the model builder intellectual satisfaction, but it is unclear whether such a model is useful in understanding a world in which agents' cognitive capacities are severely restricted.

An important shortcoming of the behavioural model presented in this article is that it does not introduce financial markets and the banking sector. Financial markets have been shown to be gripped by movements of optimism and pessimism leading to bubbles and crashes. It will be interesting to extend the model to incorporate these features and to see how they interact with the animal spirits analysed here.

References

Akerlof, G. and R. Shiller (2009), *Animal Spirits. How Human Psychology Drives the Economy and Why It Matters for Global Capitalism*, Princeton, NJ: Princeton University Press.

Anderson, S., A. de Palma and J.-F. Thisse (1992), *Discrete Choice Theory of Product Differentiation*, Cambridge, MA: MIT Press.

Binder, M. and M.H. Pesaran (1996), "Multivariate rational expectations models and macroeconomic modeling: A review and some results", in M.H. Pesaran and M. Wickens (eds), *Handbook of Applied Econometrics: Macroeconomics*, Oxford, UK: Blackwell Publishers.

Branch, W. and G. Evans (2006), "Intrinsic heterogeneity in expectation formation", *Journal of Economic Theory*, 127, pp. 264–95.

Branch, W. and G. Evans (2009), "Monetary policy with heterogeneous expectations", mimeo, Stanford University.

Brazier, A., R. Harrison, M. King, and T. Yates (2006), "The danger of inflating expectations of macroeconomic stability: Heuristic switching in an overlapping generations monetary model", Working Paper n° 303, Bank of England, August.

Brock, W. and C. Hommes (1997), "A rational route to randomness", *Econometrica*, 65, pp. 1059–95.

Christiano, L., M. Eichenbaum and C. Evans (2001), "Nominal rigidities and the dynamic effects of a shock to monetary policy", NBER Working Paper n° 8403, July.

Clarida, R., J. Gali and M. Gertler (1999), "The science of monetary policy, a new Keynesian perspective", *Journal of Economic Literature*, 37, pp. 1661–707.

De Grauwe, P. and M. Grimaldi (2006), *The Exchange Rate in a Behavioural Finance Framework*, Princeton, NJ: Princeton University Press.

Evans, G. and S. Honkapohja (2001), *Learning and Expectations in Macroeconomics*, Princeton, NJ: Princeton University Press.

Gaspar, V., F. Smets and D. Vestin (2006), "Adaptive Learning, Persistence and Optimal Monetary Policy", Working Paper Series, n° 644, European Central Bank.

Howitt, P., (2008), "Macroeconomics with intelligent autonomous agents", in R. Farmer (ed.), *Macroeconomics in the Small and the Large: Essays on*

Microfoundations, Macroeconomic Applications and Economic History in Honor of Axel Leijonhufvud, Cheltenham, UK, and Northampton, MA, USA: Edward Elgar.

Kahneman, D. and A. Tversky (1973), "Prospect theory: An analysis of decisions under risk", *Econometrica*, 47, pp. 313–27.

LeBaron, B. and L. Tesfatsion (2008), "Modeling macroeconomies as open-ended dynamic systems of interacting agents", *American Economic Review (Papers & Proceedings)*, 98 (2), pp. 246–50.

Milani, F. (2007), "Learning and time-varying macroeconomic volatility", mimeo, University of California, Irvine.

Orphanides, A. and J. Williams (2004), "Robust monetary policy with imperfect information", Board of Governors of the Federal Reserve System.

Sargent, T. (1993), *Bounded Rationality in Macroeconomics*, Oxford: Oxford University Press.

Smets, F. and R. Wouters (2003), "An estimated dynamic stochastic general equilibrium model", *Journal of the European Economic Association*, 1, pp. 1123–75.

Svensson, L. (1997), "Inflation forecast targeting: Implementing and monitoring inflation targets", *European Economic Review,* 41, pp. 11–46.

Tesfatsion, L. (2006), "Agent-based computational economics: A constructive approach to economic theory", in L. Tesfatsion and K. Judd (eds), *Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics*, Handbooks in Economics Series, the Netherlands: North-Holland/Elsevier, Spring, pp. 831–80.

Wieland, V. (2011), "Model comparison and robustness: A proposal for policy analysis after the financial crisis", in R.M. Solow and J.-P. Touffut (eds), *What's Right with Macroeconomics?*, Cheltenham, UK, and Northampton, MA, USA: Edward Elgar, forthcoming.

Woodford, M., (2003), "Interest and prices: Foundations of a theory of monetary policy", Princeton, NJ: Princeton University Press.

Woodford, M. (2009), "Convergence in macroeconomics: Elements of the new synthesis", *American Economic Journal: Macroeconomics*, 1 (1), January, pp. 267–97.

Appendix: Parameter values of the calibrated model

Heuristic model

pstar = 0;	% the central bank's inflation target
a1 = 0.5;	% coefficient of expected output in output equation
a2 = -0.2;	% a is the interest elasticity of output demand
b1 = 0.5;	% b1 is coefficient of expected inflation in inflation equation
b2 = 0.05;	% b2 is coefficient of output in inflation equation
c1 = 1.5;	% c1 is coefficient of inflation in Taylor equation
c2 = 0.5;	% c2 is coefficient of output in Taylor equation
c3 = 0.5;	% interest smoothing parameter in Taylor equation
β = 1;	% fixed divergence in beliefs
δ = 2;	% variable component in divergence of beliefs
gamma = 1;	% intensity of choice parameter
sigma1 = 0.5;	% standard deviation shocks output
sigma2 = 0.5;	% standard deviation shocks inflation
sigma3 = 0.5;	% standard deviation shocks Taylor
rho=0.5;	% rho measures the speed of declining weights in mean squares errors (memory parameter)

Rational model

pstar = 0;	% the central bank's inflation target
a1 = 0.5;	% coefficient of expected output in output equation
a2 = -0.2;	% a is the interest elasticity of output demand
b1 = 0.5;	% b1 is coefficient of expected inflation in inflation equation
b2 = 0.05;	% b2 is coefficient of output in inflation equation
c1 = 1.5;	% c1 is coefficient of inflation in Taylor equation
c2 = 0.5;	% c2 is coefficient of output in Taylor equation
c3 = 0.5;	% interest smoothing parameter in Taylor equation
sigma1 = 0.5;	% standard deviation shocks output
sigma2 = 0.5;	% standard deviation shocks inflation
sigma3 = 0.5;	% standard deviation shocks Taylor