Optimal Policy Business Cycles*

by

Victor Ginsburgh
Université Libre de Bruxelles and CORE

and

Philippe Michel
GREQAM, Université d'Aix-Marseille II

Abstract

The effectiveness of economic policies depends on the nature of expectations. Under adaptive expectations, the Philipps curve allows a government to "surprise" agents. Under rational expectations, there is less room for economic policies. We assume that only an (endogenously determined) proportion of agents form rational expectations and show that this leads the government to optimal policies which result in a policy cycle with real effects.

JEL classification: C6, E10, E32.

Corresponding author: Victor Ginsburgh, Université Libre de Bruxelles, C.P.139, 50, avenue F. Roosevelt, 1050 Brussels, Belgium. Email: vginsbur@ulb.ac.be

Published Journal of *Economic Dynamics and Control* 22 (1998), 503-518.

* This research was undertaken under contract PAI n°26 financed by the Belgian Government. We are grateful to Russel Cooper, Bertrand Crettez and two referees for extremely helpful comments, as well as for references which we had missed.

1. Introduction

The effectiveness of economic policies is known to depend on the nature of expectations as well as on the possibility for a government to precommit. When agents are not perfectly informed, the policy maker may intervene in an otherwise Pareto inefficient economy. Under rational expectations, economic policies usually have no effect on output.

Reality may well be halfway, in the sense that some agents may form naive expectations, while others are rational. Such a formulation has been suggested in several recent papers.¹ Here, we follow Crettez and Michel (1992), who assume that rational agents bear costs of information to become "rational," while naive agents do not. This leads to define an *equilibrium with economically rational expectations* as an equilibrium in which the choice made by every agent to acquire the information (or not) is optimal: in equilibrium, no agent can make himself better off by changing his choice.

In this paper, we interpret such an equilibrium as follows. Assume that at some date, all agents have made their choice on the basis of some rule of thumb, say adaptive expectations. When later on, the government announces its policy, all agents are becoming perfectly informed. However, since adapting the previous decision made to the new information implies some cost, only those for whom the cost is smaller than the benefit will use this information. In the economy as a whole, this is as if *some* agents only (those for whom the cost is larger than the benefit) had adaptive expectations according to which they make their decision. This seems an attractive framework when there is a lag between a decision and its implementation. Agents have, say, to make an investment decision and only learn about the true state of the economy later on. Some of them may wish to adjust their previous choice to the new environment, but this change will obviously imply costs. An alternative interpretation is that for some agents it is worthwhile to find out more carefully what the future will be, and to bear a cost for this knowledge, by paying consultants, conducting surveys, etc.

In each period, there will thus be an endogenous proportion of agents who use the new information, and this proportion will vary over time: When the change in the rate of inflation is small, few agents will change their decision, since for many, the gain will be lower than the cost. On the contrary, large changes in the rate of inflation (or deflation) generate important losses, and more agents will incur the cost of changing their previous

¹ Akerlof and Yellen (1985a, 1985b), Bomfin and Diebold (1992), Bary (1982), Evans et al. (1993), Haltiwanger and Waldman (1985, 1989), Krusell and Smith (1993).

decision. This leads to think that in unstable situations, a large proportion of agents will modify their previous choice, given the new information. In a stable economy, adaptive expectations do as well and in the limit, i.e. in the stationary state, myopic expectations turn out to coincide with rational expectations.

In a world where the natural rate of unemployment is perceived as too large, a government may be willing to reduce unemployment. Since there is room for economic policies if agents base their actions on rules of thumb (and agents will do so under small shocks), the government will take limited actions, and proceed by a sequence of small inflationary steps to reduce unemployment. Inflationary pressures will eventually feed into adaptive expectations and inflationary shocks must be made larger in order for the government to surprise agents. Since this cannot go forever and inflation rates must, at some point, return to normal, restrictive policies will be necessary. These have a negative impact on the government's objective function as well as on real output, unless agents find it optimal to use the information. This will be the case if the deflationary shock is made large enough, since then, to avoid large losses due to the surprise, agents will rationally modify their first decision. This leads the government to implement optimal policies which generate a business cycle: A sequence of small (but monotonically increasing) inflationary shocks will be followed by a large deflation, enabling again the government to resume its small inflationary shocks. In such a model, agents adapt their behaviour to that of the government. As suggested by Lucas (1981), the coefficients of the reduced form of the model are not invariant to policy changes.

Similar frameworks are considered by Akerlof and Yellen (1985a, 1985b), Bomfin and Diebold (1992), Bray (1982), Evans and Ramey (1992), Evans et al. (1993), Haltiwanger and Waldman (1985, 1989) and Krusell and Smith (1993). They all assume that agents bear transaction costs to "become rational" (or may be heterogeneous in terms of their information processing capabilities) and show that even a very small number of naive agents may have a disproportionately large effect on equilibrium outcomes. In particular, the policy ineffectiveness property induced by rational expectations may be offset. Haltiwanger and Waldman (1985, 1989) prove that under the additional assumption of strategic complements, the economy can be characterized by multipliers that are fairly close to Keynesian ones. Bomfin and Diebold (1992) show that the Keynesian effects are magnified if strategic complements are compounded with imperfect competition. Evans et al. (1993) show that if the proportion of agents with limited knowledge is large enough, deterministic fluctuations that are otherwise generated by the overlapping generations model they consider, may disappear.

In all these models, the number of naive agents is fixed. Krusell and Smith (1993) endogenize this number and show that consumers may decide to depart from intertemporal optimal choices and follow a rule of thumb (like a constant savings rate), even if information costs are very low (of the order of 0.04% of total consumption). Our paper considers a policy maker who sets the inflation rate in order to surprise naive agents, but these agents also optimally choose to be naive if the surprise is not too large. The number of naive agents is thus endogenous and changes over the cycle, in relation to the inflation rate chosen by the policy maker. The combined behaviour produces cycles with nonconstant surprise inflations, which have real effects. The model has positive rather than normative implications, since the government only looks at the activity level, without taking into account the costs borne by private agents.

The paper is organized as follows. Section 2 describes the basic framework, which is similar to the one used by Barro and Gordon (1983). Section 3 studies the dynamics of optimal government policies when the number of agents who form rational expectations is given, while in Section 4, we introduce the assumption of economically rational expectations. To compute the government's optimal policy, one has to solve a nonconvex dynamic program. This can only be done by numerical simulation, through which we illustrate the optimal policy cycle and its deformations when the main parameters of the economy are changed. In Appendix 1 however, we provide an example which can be solved analytically. Section 5 discusses the robustness of the results and Section 6 concludes the paper.

2. The model

We consider a simple classical reduced form model (à la Lucas-Barro-Gordon) in which agents form economically rational expectations. We assume that the deviation y_t of (the log of) output from its natural level y^* is given by:

$$(y_t - y^*) = \gamma(\pi_t - \pi_t^e), \gamma > 0, \quad (2.1)$$

where π_t and π_t^e represent the actual and the expected rate on inflation. Only the unexpected inflation rate has real effects.

The policy maker's objective is a linear-quadratic function, which involves, for each period, a net gain g_t , given by:

$$g_t = \alpha(y_t - y^*) - \frac{1}{2} (\pi_t)^2, \quad \alpha > 0. \quad (2.2)$$

The first term is the benefit from increases in output (or decreases in unemployment) resulting from unexpected inflationary shocks, while the second one represents the cost of inflation. Replacing (2.1) in (2.2) leads to:

$$g_t = \beta(\pi_t - \pi_t^e) - \frac{1}{2} (\pi_t)^2, \quad (2.3)$$

where $\beta = \alpha\gamma$.

The policymaker's intertemporal objective is to maximize the discounted sum of net gains:

$$\max \sum_{t=0}^{\infty} \delta^t g_t. \quad (2.4)$$

If agents form rational expectations, at equilibrium, $\pi_t = \pi_t^e$ and the optimal policy is evidently $\pi_t = 0$ for all t . This policy can be implemented assuming that the government has the possibility to precommit itself at each period t . If agents use a rule of thumb, e.g. form expectations that are assumed to be adaptive, one has:

$$\pi_{t+1}^a = \lambda \pi_t + (1-\lambda) \pi_t^a, \quad 0 < \lambda \leq 1, \quad (2.5)$$

with $\pi_t^e = \pi_t^a$.

Following Akerlof and Yellen (1985a, 1985b), Bomfield and Diebold (1992), Haltiwanger and Waldman (1985, 1989) and Krusell and Smith (1993), we now consider the intermediate case, in which a proportion x_t of agents use rational expectations, while a proportion $1-x_t$ use adaptive expectations. We assume that the average expected rate of inflation is a weighted average of the rates expected by rational and adaptive agents:

$$\pi_t^e = x_t \pi_t + (1-x_t) \pi_t^a. \quad (2.6)$$

In Section 3, we study the dynamics of the model when $x_t = x$ is constant and obtain the usual conclusions. In Section 4, the number of smart agents will be made endogenous: each agent may optimally choose to be smart or naive and the government will have the extra possibility to influence the relative proportions of smart and naive agents.

3. Policies when the number of naive agents is exogenous

When $x_t = x$, the first-order necessary conditions for this problem are:

$$\pi_t = \beta(1-x) - \delta\lambda q_{t+1}, \quad (3.1a)$$

$$q_t = \beta(1-x) + \delta(1-\lambda)q_{t+1}, \quad (3.1b)$$

where q_t is the shadow price associated with (2.5). It is easy to verify that the optimal inflation rate is constant and equal to:

$$\pi^* = \frac{\beta(1-x)(1-\delta)}{1-\delta+\delta\lambda}. \quad (3.2)$$

The corresponding dynamics of adaptive expectations are:

$$\pi_t^a = \pi^* + (\pi_0^a - \pi^*)(1-\lambda)^t, \quad (3.3)$$

where π_0^a is the value of π_t^a at $t = 0$.

The optimal inflation rate increases with the proportion $1-x$ of agents having adaptive expectations. It is maximal if all agents are naive ($x = 0$) and it is equal to zero if they are all rational ($x = 1$). It also increases with $(1-\lambda)$, the degree of inertia of adaptive expectations.

As long as $x < 1$, real effects, equal to:

$$(y_t - y^*) = \gamma(1-x)[\pi^*(1-\lambda)^t - \pi_0^a] \quad (3.4)$$

are persistent, but when the number of naive agents is small, they will be small also: near-rational expectations produce near-classical equilibria. The slope of the Philipps curve is

now $\gamma(1-x)$ instead of γ , and takes into account the proportion $(1-x)$ of naive agents. If all agents are rational, the usual ineffectiveness conclusion obtains since $y_t - y^* = 0$.

This optimal policy is time inconsistent like in the special case with pure rational expectations ($x = 1$) considered by Kydland and Prescott (1977). The same game-theoretic approach as the one suggested by Barro and Gordon (1983) can be followed to enforce the optimal solution, and make it time consistent.

4. Policies under economically rational expectations

We assume that at time t , decisions have been made at no cost on the basis of adaptive expectations π_t^a . An agent θ can modify his decision at a fixed cost c using the new information π_t .² There is a continuum of agents θ , $\theta \in [0, 1]$. Agent θ loses $\theta(\pi_t^a - \pi_t)^2$ when he uses π_t^a while the actual inflation rate is π_t . Let $\underline{\theta}_t$ be defined by:

$$\underline{\theta}_t(\pi_t^a - \pi_t)^2 = c. \quad (4.1)$$

Then, the loss of agent θ is larger than c iff $\theta > \underline{\theta}_t$ and the proportion x_t of agents who decide to change their decision is:

$$x_t = x(\pi_t^a, \pi_t) = \max \{0, 1 - c(\pi_t^a - \pi_t)^{-2}\}. \quad (4.2)$$

We consider the problem of a policy maker who maximizes (2.4) under (2.5), (2.6) and (4.2).

Combining (2.1), (2.5) and (2.6), one can express real effects as:

$$y_t - y^* = \gamma(1 - x_t)(\pi_t - \pi_t^a). \quad (4.3)$$

² Actually, decisions are made at time t_1 , and the new information becomes available in $t_2 > t_1$. To simplify somewhat the presentation (and the computations), we assume that $t_2 = t_1 = t$. An obvious alternative would be to replace (2.5) by $\pi_{t+1}^a = \lambda\pi_{t-1} + (1-\lambda)\pi_{t-1}^a$, $0 < \lambda \leq 1$.

This expression shows that real effects are directly linked to the proportion $(1 - x_t)$ of naive agents who expect an inflation rate π_t^a , while the actual rate is π_t . Replacing x_t in (4.3) by its value (4.2) leads to:

$$y_t - y^* = \gamma c (\pi_t - \pi_t^a)^{-1} \text{ if } (\pi_t - \pi_t^a)^2 \geq c,$$

or

$$y_t - y^* = \gamma (\pi_t - \pi_t^a) \text{ if } (\pi_t - \pi_t^a)^2 \leq c.$$

The maximal real gain $(y_t - y^*)$ is obtained for $(\pi_t - \pi_t^a) = c$, and is equal to:

$$y_t - y^* = \gamma \sqrt{c}. \quad (4.4)$$

This shows that for any $c > 0$, policy has real effects. As long as c and π_t^a are not too large, the cost induced by inflation in the policy maker's objective $(-\frac{1}{2} \pi_t^2)$ is dominated by the real effects terms $\beta(\pi_t - \pi_t^e)$ and it is optimal for the policy maker to increase the inflation rate.

The solution corresponds to an intertemporal equilibrium under economically rational expectations (in the sense defined by Crettez and Michel (1992)), in which every agent makes an optimal choice. Under the alternative interpretation discussed in the introduction, the solution is an intertemporal equilibrium in which, ex ante, agents use adaptive expectations to make their choices. Ex post, they all know the exact value of the inflation rate, but only those for which a change of decision is beneficial make a move.

Solving the model

The policy maker problem (2.1)-(2.4) with x_t defined by (3.2) is neither concave nor differentiable and can thus not be analyzed by the usual optimization tools (calculus of variations or Pontryagin's maximum principle). It can however be solved using dynamic programming.³

In Appendix 1, we give a simple case in which (a) inflation rates can take only three values and (b) agents bear no cost when their expectation error is small and infinitely

³ Bellman's principle gives the TC1 equilibrium (see Cohen and Michel, 1988). Here, this equilibrium is the first best, since the agents' behaviour depends on their expectations of the current government policy only.

large costs otherwise. This is sufficient to exhibit an optimal policy cycle in which two low inflation periods (during which no agent modifies his decision) are followed by a large deflation (in which all agents use the new information).

Less trivial cases need to be examined by numerical simulation; we solve Bellman's equation by computing the value function and the solution on a grid of points.⁴ The value function of the dynamic program can be written:

$$V(\pi_0^a) = \max \{ \Sigma \delta^t \phi(\pi_t^a, \pi_t) \text{ subject to (2.3)} \}, \quad (4.5)$$

where

$$\begin{aligned} \phi(\pi_t^a, \pi_t) &= \beta(\pi_t - \pi_t^e) - \frac{1}{2} \pi_t^2, \\ \pi_t^e &= x_t \pi_t + (1 - x_t) \pi_t^a, \\ x_t &= \max \{ 0, 1 - c(\pi_t^a - \pi_t)^{-2} \}. \end{aligned}$$

Bellman's equation for this problem is then:

$$V(\pi^a) = \max_{\pi} \phi(\pi^a, \pi) + \delta V(\lambda \pi + (1-\lambda)\pi^a). \quad (4.6)$$

We chose parameter values in order to obtain reasonable values for real output gains, generated by reasonable changes in inflation rates. From (2.3), it can be seen that the maximal gain which can be generated by the policy maker is equal to β (this is the case of a full surprise with $\pi_t^e = \pi_t^a$). We chose this maximal rate to be 10% and this implies setting $\beta = 0.10$. From (4.4), this maximal rate is also equal to $\gamma \sqrt{c}$. Setting $\gamma = 1$, it follows that $\sqrt{c} = 0.01$ and, therefore, $c = 0.0001$. Note also that this is the largest loss incurred by agents when the error they make in anticipating inflation is equal to 1% (equation (4.1) with $\theta = 1$). Finally, we set $\lambda = 0.75$ and $\delta = 0.95$. These values generate optimal inflation rates varying between -0.024 and $.017$ and real effects on output varying between -0.002 and $.010$. The basic result is illustrated in Figure 1.

Simulating the model

As expected, for the case of zero information costs ($c = 0$), the algorithm finds the solution of Section 3, with $x_t = 1$, since all agents use all the available information without

⁴ The algorithm used to compute the solution V of this equation is briefly described in Appendix 2.

cost. For small but positive values of c , say $c = 10^{-10}$, we obtain a rather long cycle, with very small amplitude and negligible real effects. For smaller values of c , the length of the cycle increases, while its amplitude decreases and in the limit, for c very small, the cycle may be considered to be of infinite length, converging to the situation with $c = 0$.

When c increases, the length of the cycle decreases and its amplitude increases, as is illustrated in Figure 2. This behaviour is characteristic of the optimal policy cycle. At the end of the cycle, there is more and more inflation. This imposes large costs to the government (term $-\frac{1}{2}(\pi_t)^2$ in the objective function). To restart the cycle with lower inflation rates, the government has to generate a deflationary shock, which will cost him $\beta(\pi_t - \pi_t^e)$. Since this shock is taken into account by agents, the cost will be small. To make sure that as many agents as possible modify their decision, the government has to generate a large shock, so that, in equilibrium, agents will find it more profitable to spend c and use all the information. Hence, the larger the cost, the larger the optimal deflationary shock.

The shortest cycle has three periods (two inflationary ones, and a unique deflation) and is obtained for $c = .01$. When c increases further, cycles become longer. Their typical shape can be described as follows. There is a series of small inflations, followed by a constant rate that corresponds to the long term optimal inflation rate π^* , and ending with a large deflation. This shape is illustrated in Figures 3 and 4 for two different values of the cost parameter. The model behaves like a linear quadratic model during the part of the cycle in which the policy maker chooses small inflationary steps (numerically, we find that the number of rational agents is zero during that phase) and when the inflation steps are constant, adaptive expectations give the same result as rational expectations: Agents have time to learn.

Cycles eventually tend to vanish (or become infinitely long) for some value c° (given our calibration, $c^\circ = .01462$). The reason is the following. The existence of the policy cycle is due to the possibility for the government to decide on a large deflation, which will cost little in terms of output, since agents adjust to this by changing their decisions (i.e. by becoming rational). However, when the cost becomes too large, agents get more reluctant to adjust to small shocks; therefore, these have to become larger. However, even if the deflation is perfectly anticipated and has no effect on real output, it bears heavily on the government objective (2.2). It so happens that for $c \geq c^\circ$, the cost to the government of a deflation is such that it will only generate deflations which are too

small to induce enough agents to become rational and to spare the economy from a drop in real output.

It is obviously the cost c which is the important parameter. If this cost is large, it is not optimal for the government to generate cycles, since the deflation which necessarily ends a cycle will generate costs that supersede the gains obtained during inflationary steps.

Simulations in which the values of other parameters are modified, show that:

- (a) an increase in β , the weight on the benefit to the government, resulting from unexpected inflation shocks, increases the average optimal inflation rate, without affecting the shape of the cycle;
- (b) a decrease in the discount factor δ , increases the amplitude of the cycle;⁵
- (c) a decrease in the adjustment speed λ increases both the amplitude and the length of the cycle.

5. Discussion of the results

The simple model described above leads to results that look extreme. However, introducing economically rational expectations has seriously changed the conclusions reached by a Barro-Gordon type model. Constant economic policies are replaced by cycles, that are necessarily regular since there are no exogenous (or random) shocks, agents do not learn from the past, and expectations are formed according to systematic rules.

Adding random shocks would considerably complicate the problem solved by the government. It is very likely that the model would also produce cycles, though these would be irregular. It would then be interesting to check whether, under a reasonable parametrization, the model produces cycles that have the same moments as observed cycles.

⁵ Note that in the example studied in Appendix 1, a cycle will exist for every value of δ .

If agents knew the government's objective (maybe after some learning), we would be back in the model with rational expectations and no information costs, so that the policy would be constant, like in the Barro-Gordon model.

In our formulation of economically rational expectations, agents are either rational or form adaptive expectations, but other forms of naive expectations could be envisaged. The important point is that the proportions of both types of agents is endogenous and depends on the policy decisions. It is easy to imagine other kinds of economically rational expectations that would not lead to cycles. However, with enough regularity in the formation of non-rational expectations, we conjecture that cycles would persist, since the association "low inflation naive expectations surprises" and "high deflation rational expectations no surprises" would still drive the results.

Finally, instead of considering a reduced-form model, it would be possible to describe the behaviour of agents in a more explicit way. However, endogenizing the proportion of agents who behave rationally is likely to lead to extra technical difficulties that we wanted to avoid at this stage, by using the simplest possible framework.

6. Conclusion

We have used a very simple model to show that the assumption of economically rational expectations may lead to qualitative results which are very different from those generated by standard assumptions on expectations. The resulting cycles are obtained by combining two polar assumptions on the formation of expectations. Clearly, if it were true that the government generated regular cycles, then agents would have learnt from the past and cycles would disappear (the Barro-Gordon framework). This would of course not be so if the economy were subject to random disturbances. Nevertheless, as soon as one generalizes the idea of rational expectations and takes into account information costs, the relation between economic policy and expectations is modified. This had already been pointed out in other papers, discussing models in which some agents did not behave rationally.

The novelty here is that the number of agents who behave rationally is endogenous and depends on the action of the government. Small inflations are "ignored" (agents behave naively) and have therefore effects on real output, whereas all agents react to important deflations by becoming rational and this has little, if any, real effects. This leads

the government to generate policy cycles which consist in several inflationary steps (which increase real output with respect to its natural level), followed by a deflation (with almost no decrease in output). The model thus shows that disinflation can be achieved at low cost, provided that the government avoids gradualism, since large shocks are perfectly anticipated and do not affect real output too much. This conclusion differs substantially from the results obtained in competitive equilibrium cycles discussed in optimal growth theory (see e.g. Benhabib and Nishimura (1985)).

References

- Akerlof, G.A. and J.L. Yellen, 1985a, A near-rational model of the business cycle, with wage and price inertia, *Quarterly Journal of Economics* 50, 823-838.
- Akerlof, G.A. and J.L. Yellen, 1985b, Can small deviations from rationality make significant differences to economic equilibria, *American Economic Review* 75, 708-720.
- Barro, R.J. and D.B. Gordon, 1983, Rules, discretion and reputation in a model of monetary policy, *Journal of Monetary Economics* 12, 101-121.
- Bellman, R., 1961, *Adaptive Control Processes* (Princeton University Press).
- Benhabib, J. and K. Nishimura, 1985, Competitive equilibrium cycles, *Journal of Economic Theory* 35, 284-306.
- Bomfin, A.N. and F.X. Diebold, 1992, Near-rationality and strategic complementarity in a macroeconomic model: policy effects, persistence and multipliers, Discussion Paper, Federal Reserve Board, Washington, D.C.
- Bray, M., 1982, Learning, estimation and the stability of rational expectations equilibria, *Journal of Economic Theory* 26, 318-339.
- Cohen, D. and Ph. Michel, 1988, How should control theory be used to calculate a time-consistent government policy, *Review of Economic Studies* LV, 263-274.
- Crettez, B. and Ph. Michel, 1992, Economically rational expectations equilibrium, *Economics Letters* 40, 203-206.
- Evans, G.W., S. Honkapohja and T.J. Sargent, 1993, On the preservation of deterministic cycles when some agents perceive them to be random fluctuations, *Journal of Economic Dynamics and Control* 17, 705-721.
- Evans, G.W. and G. Ramey, 1992, Expectation calculation and macroeconomic dynamics, *American Economic Review* 82, 207-223.
- Haltiwanger, J. and M. Waldman, 1985, Rational expectations and the limits of rationality: an analysis of heterogeneity, *American Economic Review* 75, 326-340.
- Haltiwanger, J. and M. Waldman, 1989, Limited rationality and strategic complements: the implications for macroeconomics, *Quarterly Journal of Economics* 104, 463-483.
- Krusell, P. and A.A. Smith, Jr., 1993, Costly implementation and equilibrium decision rule selection in the stochastic growth model, Discussion Paper, Carnegie Mellon University, Pittsburgh, PA.
- Kydland, F.E. and E.W. Prescott, 1977 Rules rather than discretion: the inconsistency of optimal plans, *Journal of Political Economy* 85, 473-491.
- Lucas, R., 1981, Econometric policy evaluation: a critique, in R. Lucas, *Studies in Business-Cycle Theory* (MIT Press, Cambridge, MA)..

Figure 1
Inflation and real output cycles for the base case
 ($\beta = 0.1$, $c = 0.0001$)

Figure 2
Real output cycles
for various information costs
 (values x 1,000)

Figure 3
Inflation for large information costs

Figure 4
Real output for large information costs

Appendix 1. A simple case exhibiting an optimal cycle

The net gain for the government is given by the following parametrization of (2.1):

$$g_t = 2(\pi_t - \pi_t^e) - \frac{1}{2}(\pi_t)^2.$$

Expectations are naive, so that (2.3) reads:

$$\pi_{t+1}^a = \pi_t. \quad (\text{A1})$$

All agents are identical and lose 0 when the forecasting error is not larger than 1 (i.e. when $|\pi_t^a - \pi_t| \leq 1$); they bear an infinitely large cost when the error is larger than 1; it follows that at equilibrium, the proportion x_t of agents buying perfect forecasts (eq. (3.2)) is given by:

$$\begin{aligned} x_t = x(\pi_t^a, \pi_t) &= 0 \text{ if } |\pi_t^a - \pi_t| \leq 1, \\ x_t = x(\pi_t^a, \pi_t) &= 1 \text{ if } |\pi_t^a - \pi_t| > 1. \end{aligned} \quad (\text{A2})$$

The expected rate of inflation is given by (2.4):

$$\pi_t^e = x_t \pi_t + (1-x_t) \pi_t^a. \quad (\text{A3})$$

For simplicity, we assume that there can be only three values for the inflation rate π_t :

$$\pi_t \in \{0, 1, 2\}.$$

We consider the following policy rule:

$$\begin{aligned} \pi^*(\pi^a) &= 1 \text{ if } \pi^a = 0 \\ \pi^*(\pi^a) &= 2 \text{ if } \pi^a = 1 \\ \pi^*(\pi^a) &= 0 \text{ if } \pi^a = 2. \end{aligned} \quad (\text{A4})$$

We first show that this rule generates a policy cycle. Suppose indeed that we start in period t with $\pi_t^a = 0$. Then, by (A4), $\pi_t^* = \pi^*(\pi_t^a) = 1$; by (A2), $x_t = 0$ and by (A3), $\pi_t^e = 0$. For period $t+1$, $\pi_{t+1}^a = \pi_t^* = 1$, $\pi_{t+1}^* = \pi^*(\pi_{t+1}^a) = 2$, $x_{t+1} = 0$ and $\pi_{t+1}^e = 1$. Finally, in $t+2$, $\pi_{t+2}^a = \pi_{t+1}^* = 2$, $\pi_{t+2}^* = \pi^*(\pi_{t+2}^a) = 0$, $x_{t+2} = 1$ and $\pi_{t+2}^e = 0$. It is easy to see that in period $t+3$, we are back with $\pi_{t+3}^a = \pi_{t+2}^* = 0$ and that the same cycle restarts.

In order to show that this policy rule is indeed optimal, we shall use the value function defined by:

$$V(\pi^a) = \max_{\pi} \phi(\pi^a, \pi) + \delta V(\pi),$$

where $\phi(\pi^a, \pi) = g(\pi, \pi^e) = g(\pi, x\pi + (1-x)\pi^a)$ is the instantaneous gain of the government. For further use, we compute the values of this function for values of $\pi \in \{0, 1, 2\}$ and $\pi^a \in \{0, 1, 2\}$ with $x = x(\pi^a, \pi)$ given by (A2):

		π		
		0	1	2
π^a	0	0	3/2	-2
	1	-2	-1/2	0
	2	0	-5/2	-2

If rule (A4) is optimal, the value function must verify:

$$\begin{aligned} V(0) &= \phi(0, \pi^*) + \delta V(\pi^*), \text{ with } \pi^* = \pi^*(0) = 1 \\ &= \phi(0, 1) + \delta V(1) = 3/2 + \delta V(1), \\ V(1) &= \phi(1, \pi^*) + \delta V(\pi^*), \text{ with } \pi^* = \pi^*(1) = 2 \\ &= \phi(1, 2) + \delta V(2) = \delta V(2), \\ V(2) &= \phi(2, \pi^*) + \delta V(\pi^*), \text{ with } \pi^* = \pi^*(2) = 0 \\ &= \phi(2, 0) + \delta V(0) = \delta V(0). \end{aligned}$$

This leads to the following system of three equations:

$$V(0) = 3/2 + \delta V(1)$$

$$V(1) = \delta V(2)$$

$$V(2) = \delta V(0),$$

the solution of which is:

$$V(0) = \frac{3}{2(1-\delta^3)}; V(1) = \delta^2 V(0); V(2) = \delta V(0). \quad (\text{A5})$$

It is necessary (and sufficient) that the function defined in (A5) satisfies Bellman's equation:

$$V(\pi^a) = \max_{\pi} \phi(\pi^a, \pi) + \delta V(\pi)$$

for every value of π^a ; indeed,

- for $\pi^a = 0$: $\max \{ \delta V(0), 3/2 + \delta V(1), -2 + \delta V(2) \} = V(0)$,
- for $\pi^a = 1$: $\max \{ -2 + \delta V(0), -1/2 + \delta V(1), \delta V(2) \} = V(1)$,
- for $\pi^a = 2$: $\max \{ \delta V(0), -5/2 + \delta V(1), -2 + \delta V(2) \} = V(2)$,

which shows that rule (A4) defined above is optimal.

Appendix 2. Computing a solution to equation (4.6)

In order to compute a solution to the functional equation (4.6), one constructs the following sequence:

$$V_{k+1}(\pi^a) = \max_{\pi} \phi(\pi^a, \pi) + \delta V_k(\lambda\pi + (1-\lambda)\pi^a), \quad (\text{A6})$$

with $V_0(\pi^a)$ chosen to be equal to zero; the iterates are stopped when V_k is (approximately) equal to V_{k-1} .

Numerically, one sets an interval of variation for π^a and π , say $[a, b]$ and one chooses a number of points N on each interval. The setting of the interval is delicate: we start with $[a, b] = [-0.15, 0.10]$ and when this interval shows to be too large, given the values of the solution, we gradually narrow it down to increase precision (the number of grid points is kept constant, so that precision is increased when the interval gets smaller). For each point in $[a, b]^2$, one computes $\phi(\pi^a, \pi)$ and $f = \lambda\pi + (1-\lambda)\pi^a$. To correct for numerical errors due to the discrete grid on $[a, b]^2$, $V_k(f)$ in (A6) is replaced by the interpolation $\mu V_k(f_-) + (1-\mu)V_k(f_+)$, where f_- is a point in the $[a, b]$ interval for π^a which precedes f and f_+ is a point which follows with $\mu = f - f_-$. See Bellman (1961, chapter 5) for details.

The optimal trajectory for inflation, starting at π_0^a is obtained by:

$$\pi_{t+1}^a = \lambda \pi_t^*(\pi_t^a) + (1-\lambda)\pi_t^a,$$

where $\pi_t^*(\pi_t^a)$ is the optimal policy rule, i.e. is the solution π^* that maximizes (4.6).