



Discussion Paper Series

Escape Dynamics and Policy Specification

By

Michele Berardi

Centre for Growth and Business Cycle Research, Economic Studies,
University of Manchester, Manchester, M13 9PL, UK

March 2009
Number 117

Download paper from:

<http://www.socialsciences.manchester.ac.uk/cgbcr/discussionpapers/index.html>

Escape Dynamics and Policy Specification

Michele Berardi*
University of Manchester

March 2009

Abstract

In his monograph *The conquest of American inflation* (1999) Sargent suggests that the sharp reduction in US inflation that took place under Volker may vindicate the type of econometric policy evaluation famously criticized by Lucas (1976). At the core of this vindication story are the escape dynamics, recurrent sliding away from the path leading to the time-consistent sub-optimal equilibrium level of inflation. We try to understand here under which conditions this phenomenon arises. In particular, we note that economists, and consequently policymakers, knew long before the Lucas critique that in order to do policy analysis structural models were required. We thus endow our policymaker with a correctly specified model, one that takes explicitly into account the role of expectations. Using such a model, together with a policy that takes expectations as given, the escape dynamics do not appear. But they reappear when long run considerations of policy effects enter into the picture. We thus conclude that what really matters is the way in which the policymaker designs its policy, rather than the econometric specification of the model he uses.

Key words: escape dynamics, policy rule, expectations, learning.

JEL classification: E52, E58, E61.

1 Introduction

The rise of inflation that took place in the United States in the late 1960s and 70s and its subsequent sharp reduction under Volker have attracted the

*The first draft of this paper was written while I was visiting student at the Carnegie Mellon University, whose hospitality is gratefully acknowledged. I would also like to thank, without implicating, Ben McCallum and Daniele Cohen-Pirani for helpful comments and discussions at the time. All errors are my own.

attention of many economists over time and different alternative explanations have been put forward,¹ from price shocks to changes in the natural rate of unemployment.² A more convincing account of these historical events is the one that establishes a close relation between different historical inflation records and different (indeed, almost opposite) theories that happened to prevail at the time. The rise in inflation that took place starting from the late 60s was, according to this argument, the result of the view, first developed by economists and then spread out among policymakers, that a permanent trade-off between inflation and unemployment was built into the economic system and was, moreover, exploitable. Trying to pursue lower unemployment, the policymaker ended up in generating a burst in inflation. The subsequent stabilization was the result of the development by economists of the Natural Rate Hypothesis (NRH) and of its understanding by policymakers. This explanation is comforting, since it implies that those events can not happen again: the policymaker has understood the lesson, and is no more prone to an inflationary bias.

Sargent (1999) presents instead a different story, one that vindicates the type of econometric policy evaluation famously criticized by Lucas (1976). According to this story, starting from the late 1960s, and following leading economists like Samuelson, Solow and Tobin, the policymaker was inclined to believe in the existence of an exploitable trade-off between inflation and unemployment. To gain information about the response of the economy to monetary policy, the authority recurrently estimated a distributed lag Phillips curve, ignoring the warning made by Friedman and Lucas about the need to identify private sector's expectations as a relevant variable in positioning (and shifting) the Phillips curve. Fearing parameter drift, the policymaker ran his regressions discounting past observation: this practice generated estimations that fluctuated over time, validating the (wrong) initial belief of parameter drifting. These oscillations in estimated parameter values, in turn, were reflected in oscillations in the level of inflation, as the policymaker used the most recent estimate of his model to set the policy. When the estimated Phillips curve showed a positive trade-off, the authority increased inflation, in order to reduce unemployment. This practice affected the relation between macroeconomic variables, so that successive estimations revealed an adverse shift in the Phillips curve. On the basis of the new information, the policymaker perceived a less favorable trade-off and was in-

¹For a brief review of possible alternative explanations, see Taylor (1996).

²This hypothesis has been tested econometrically by Ireland (1999), finding that data on US inflation and unemployment from 1960 to 1997 are consistent with it.

duced to reduce inflation. This dynamic process, Sargent maintains, would account for the burst in inflation during the late 1960s and 70s, and for its sharp reduction during the 1980s. But it leaves open the possibility of the same events happening again, as soon as data will deceive the policymaker by showing a new exploitable trade-off.

The idea that an adaptive policy can temporarily mitigate the inflationary bias first described by Kydland and Prescott (1977) comes from Sims (1988): a policymaker who conducts his policy allowing for uncertainty and tries to exploit a perceived trade-off can learn an approximated version of the NRH and generate nearly optimal outcomes even without explicitly recognizing the role of expectations. Since policy-created variation in inflation has no effects at all on unemployment, a policymaker who periodically reconsiders the statistical relation between these two variables during its policy experiment realizes this lack of correlation and thus reduces inflation and implements a near optimal policy. This better outcome nevertheless can not persist: while low inflation is maintained, spurious statistical evidence emerges that leads the policymaker to believe again in an exploitable trade-off and consequently to increase inflation.

The purpose of this paper is to reconsider this possible alternative explanation to the historical path of U.S. inflation, and the conditions under which its main feature, the escape dynamics, manifest themselves. Economists, and consequently policymakers, knew long before the Lucas critique that in order to do policy analysis structural models were required.³ But structural models alone are not enough: we show that, using such models, together with a policy rule derived taking expectations as given, the escape dynamics that Sargent presents as the central element of the vindication story do not appear and the economy displays a persistently high inflation rate. Escape dynamics, and the possibility of (temporarily) low inflation, reappear, though, when long-run considerations are allowed to play a role in the policy decisions of the authority. We thus claim that what really matters for the outcome in terms of inflation dynamics is the way in which the policymaker derives the optimal policy, rather than the econometric specification of the model he uses.

³The need for structural models in policy analysis was recognized at least since the early 1950s. See Marschak (1953).

2 Escape dynamics, inflationary bias and optimal policy

The adaptive expectations story in Sargent (1999) departs from the rational expectations paradigm and considers how the process agents (in this case, the policymaker) follow in order to form their beliefs about the relevant economic variables can influence the outcome of an economic system. Stepping back from the equilibrium condition that lies behind the rational expectations assumption, this approach considers situations in which expectations have not yet converged to the fixed point that characterizes RE equilibria.

In Sargent (1999) set up, the policymaker tries to manage the economy by setting inflation (treated as the policy instrument) at the level that minimizes the expected losses coming from the resulting combinations of unemployment and inflation. In order to accomplish this task, the policymaker needs to know how the economy will respond to his policy, i.e., he needs to know the actual Phillips curve. Not being endowed a priori with this knowledge, the central bank (CB) learns it from experience, recurrently estimating from the data the parameters in a misspecified model,⁴ a distributed-lag Phillips curve.

If the policymaker uses a decreasing gain algorithm, the ultimate outcome is the Nash, self-confirming, time-consistent equilibrium, with the unemployment at its natural level and inflation higher than at the optimal level. An analysis of the convergence of this learning rule, using the stochastic approximation approach, is worked out by Evans and Honkapohja (2001), in ch. 13.2. Since in equilibrium both the true and the misspecified model give the same prescription in terms of optimal policy to implement, the policymaker is not disappointed by the outcome he observes even though his beliefs about the economy are wrong.

But if the government at each period re-estimates a distributed-lag Phillips curve using a constant gain algorithm, an interesting result is found: the policymaker is able to learn from past attempts to exploit a perceived trade-off between inflation and unemployment, and at times discovers a version of the natural-rate hypothesis that periodically makes him reduce inflation and lead the economy towards the Ramsey, optimal, time-inconsistent outcome (which is not an equilibrium).⁵ Sargent calls this sliding away from the path leading to the Nash equilibrium the escape route.

⁴The model is misspecified in the sense that it fails to take explicitly into account the role of expectations in positioning the Phillips curve.

⁵The Ramsey outcome is the one that would be chosen by a central planner maximizing social welfare.

Theory of large deviations, as explained in Cho, Williams and Sargent (2002), characterizes the escape dynamics. These are spurred by unusual shocks affecting the learning of the policymaker, who temporarily gets to learn an approximated version of the NRH and overcomes its inflationary bias. An heuristic analysis of the convergence can be found in Evans and Honkapohja (2001), ch. 14.4.⁶

To give a better perception of this phenomenon, we reproduce here the dynamics of the economy under the adaptive scheme with constant gain, in the simple "static" case in which the perceived Phillips curve doesn't contain any lags.⁷ This means that the policymaker solves at any time t a one-period problem, since today's policy is not perceived to influence future outcomes.⁸ In the basic setting we employ here, the true data generating process is represented by an expectation-augmented Phillips curve that embodies the NRH. as the public is supposed to have rational expectations, no pattern of monetary policy affects unemployment, which fluctuates stochastically around its natural rate level as the economy is hit by uncorrelated shocks.

In particular, the actual Phillips curve is of the type

$$U_t = U^* + \phi(\pi_t - E_{t-1}\pi_t) + \nu_{1,t}, \quad (1)$$

while the perceived Phillips curve in this static case has the form

$$U_t = a_{0,t-1} - a_{1,t-1}\pi_t + \epsilon_{1,t}. \quad (2)$$

The resulting reduced form for unemployment is

$$U_t = U^* - \phi\epsilon_{2,t} + \nu_{1,t} \quad (3)$$

where $\epsilon_{2,t}$ is the shock in the inflation equation

$$\pi_t = \hat{\pi}_t + \epsilon_{2,t} \quad (4)$$

⁶They note that the phenomenon of escape dynamics seems to be sensitive to the choice of the learning scheme: simulations using stochastic gradient learning with constant gain did not show this phenomenon. The stochastic gradient learning scheme differs from recursive least square (RLS), and in general from Newton-type algorithms, because it does not take into account the matrix of second moments in updating beliefs. We tried some simulations using our baseline model, modified to implement a stochastic gradient algorithm, and no escape dynamics showed up. This aspect, we think, deserves further attention.

⁷The same basic version of the model has been used in Gerali and Lippi (2002) and Cho, Williams and Sargent (2001).

⁸Here our policy-maker utilizes a certainty-equivalent rule, disregarding uncertainty about parameter estimates.

and $\nu_{1,t}$ is the shock in the actual Phillips curve (1)

In this model, the central bank can set the inflation rate up to a random component, as expressed in equation (4), where $\hat{\pi}_t$ is the inflation decided by the policymaker. $\hat{\pi}_t$ is therefore the policy instrument, that is chosen in order to minimize the expected loss coming from the resulting combination of unemployment and inflation, according to the equation^{9,10}

$$\min_{\hat{\pi}} E_{t-1} \sum_{t=1}^{\infty} \beta^{t-1} (U_t^2 + \pi_t^2). \quad (5)$$

The resulting policy rule¹¹ is

$$\hat{\pi}_t = \frac{a_{0,t-1} a_{1,t-1}}{1 + a_{1,t-1}^2} \quad (6)$$

Parameters in (2) are updated at any end of period t (or beginning of period $t + 1$, before deciding the new policy) when new data is available, according to the algorithm:

$$a_t = a_{t-1} + \alpha_t R_t^{-1} X_t (U_t - \hat{U}_t) \quad (7)$$

$$R_t = R_{t-1} + \alpha_t (X_t X_t' - R_{t-1}) \quad (8)$$

where

$$a_t = [a_{0,t} \quad a_{1,t}]' \quad (9)$$

$$X_t = [1 \quad \pi_t]' \quad (10)$$

$$\hat{U}_t = a_{t-1}' X_t \quad (11)$$

When α_t is a (small) constant we have the constant gain algorithm. To obtain the RLS, simply substitute α_t with $1/t$ in the learning scheme. The

⁹Here the target levels for unemployment and inflation are set to zero. Since the natural rate of unemployment in this model is 5%, the policymaker has an inflationary bias.

¹⁰As in Sargent's work, the policymaker disregards uncertainty about parameter values and utilizes the most recent estimates as if they were the true ones. In other words, the authority behaves as if the perceived law of motion was time invariant, and under this assumption derives the optimal policy. Beck and Wieland (2002) call the rule coming from this procedure the *certainty-equivalent rule*. Kendrick (2002) reports different names for this method: *sequential certainty equivalence*, *update certainty equivalence* and *heuristic certainty equivalence*.

¹¹For the derivation of this policy, see the Appendix, Case 1.

matrix R_t in (8) weighs the different elements of the vector conveying new information, giving more importance to those that are less volatile.¹²

The Phillips curve (1) emphasizes that only unexpected inflation can affect the real economy (employment or output). Since agents here are assumed to have RE, any systematic component of the policy has no real effects, even in the short run. Unemployment just fluctuates randomly around its natural level.

As the simulations show (Figure 1), inflation recurrently slides away from the path leading to the Nash equilibrium (in the simulation, this corresponds to a 5% level of inflation and a 5% level of unemployment) and approaches the Ramsey outcome (inflation at 0, unemployment at 5): when the estimated parameters make the policymaker believe in the NRH, it becomes convenient to disinflate the economy. This happens when the parameter a_1 in (2) goes to 0 (see Figure 2). But as soon as the inflation is driven towards zero, the estimated Phillips curve shows a new exploitable trade-off, and the policymaker starts trying to exploit it again.¹³

Sargent considers the more general case in which the perceived Phillips curve has lagged inflation and unemployment in it. In particular, the one utilized in his simulations is of the type

$$U_t = \gamma_{0,t-1} + \gamma_{1,t-1}\pi_t + \gamma_{2,t-1}\pi_{t-1} + \gamma_{3,t-1}\pi_{t-2} + \gamma_{4,t-1}U_{t-1} + \gamma_{5,t-1}U_{t-2} + \epsilon_t. \quad (12)$$

His results show that adding lags to the perceived Phillips curve does not change this result: periodically the policymaker finds it convenient to drive the economy towards the Ramsey, time-inconsistent output, and then to inflate it again to exploit a perceived trade-off. Again, these recurrent switches between mean dynamics, that drive the economy towards the Nash equilibrium, and escape dynamics, that point towards the Ramsey outcome, do not appear when a decreasing gain, such as $1/t$, is employed. In this case the economy converges towards the Nash equilibrium.

The alternating of these two types of dynamics, mean and escape, Sargent argues, may explain the high inflation during the 70s and the successive

¹²The initial values for a_0 and R_0 must be set. In the example, we set $a_{0,0} = 5$, $a_{1,0} = 1$ and R_0 equal to the identity matrix (we assign equal weight to the different components, supposing that at the beginning the central bank has no information about their volatility). In addition, ϕ is set equal to 1, $U^* = 5$, $\alpha = 0.0275$. The standard deviation of all the shocks is set at 0.3.

¹³Gerali and Lippi (2002) study the effects on the escape dynamics of different policy objectives by the policymaker. They find, somewhat counterintuitively, that a sufficiently inflation-averse policymaker never escapes Nash inflation. Nevertheless, his conservatism does produce a lower inflation rate on average.

reduction in the 80s in the U.S.. Using the type of econometric policy evaluation criticized by Lucas, the Federal Reserve in the late 60s and 70s found in the data a positive trade-off to be exploited, and thus inflated the economy; then, during the 80s, the induction hypothesis in the estimated Phillips curve was satisfied,¹⁴ no trade-off was perceived to exist and thus the Fed reduced inflation. But if this is the case, and the "conquest" of inflation is not due instead to a more sound understanding of the NRH by the policymaker, we are doomed to experience inflation again, as soon as the perceived Phillips curve will start showing a new exploitable trade-off.

3 Inflationary bias and policymaker's beliefs

In the case of a static Phillips curve, as represented by eq. (2), intertemporal links enter only through the learning activity. But when we introduce lagged inflation in the perceived Phillips curve (either directly or through expectations), the policymaker must take into account the direct effect that today's inflation will have (or better, is perceived to have) on tomorrow's unemployment, even if he is employing a certainty equivalent rule: the problem becomes dynamic. For any lagged term in the perceived Phillips curve there will be a forward-looking term in the policy rule derived from the optimization problem. These forward-looking terms reflect the fact that since today's unemployment depends on past inflation through the lagged component in the Phillips curve, next period's unemployment will depend on today's inflation, i.e., current policy. And since next period's unemployment will depend on next period's policy as well, current and future policies are mutually dependent. This effect must therefore be taken into account in the intertemporal optimization problem. It follows that when we derive the optimality conditions, we find that current policy depends on expected future policy. To actually implement the policy, then, this expectational term must be solved out. We do it by deriving the minimum state variable

¹⁴The induction hypothesis in the "classical" direction of fit here considered is fulfilled when the sum of coefficients on current and lag inflation is estimated to be zero. We refer to the "classical" direction of fit, as opposed to the Keynesian one, following Sargent's terminology. King and Watson (1994) attribute the same "classical" scheme to the monetarist school. The difference with the Keynesian Phillips curve is that the latter has inflation on the LHS and unemployment and lagged inflation on the RHS. In this case, the induction hypothesis is satisfied when the sum of coefficients on lagged inflation is equal to one. McCallum (1994) argues that both monetarists and keynesians used the same specification, the one with unemployment on the RHS, and that this practice didn't amount to considering this variable as exogenous, since the problem of simultaneity was acknowledged.

(MSV) solution of the model, using the undetermined coefficient procedure.

Consider the policy problem presented in the Appendix, Case 3: the MSV solution applied under the hypothesis that the true parameter values are known gives interesting results. The true Phillips curve, represented by (1), implies $a_1 = a_2 = 1$. Under this condition, the discount factor β and the parameter b_1 governing the speed of adaptation of expectations to actual inflation become crucial: with $b_1 = 1$ and $\beta = 1$ the policymaker sets inflation to zero at every period.¹⁵ The optimal, time inconsistent result is thus obtained: the policymaker follows a rule that prescribes to set inflation to zero at any period, since this is the best outcome obtainable on average, over a long period of time. In any period the policymaker could obtain better results by generating inflation, but this behavior would lead to a worst outcome in the long run. This is the same spirit of the argument advanced by McCallum (1997): on average, better results are obtained by sticking to a rule that prescribes to set inflation to zero. The policymaker recognizes this point, and he is not tempted to generate surprise inflation. At the same time, private agents understand that the authority is consistently following a rule, and expect them to keep doing it, thus solving the commitment problem.

But when β is smaller than one, the rule tells the policymaker to set a positive level of inflation. With $\beta = 0.9$ the policy prescribes a level of inflation equal to 0.5, when $\beta = 0.5$ optimal inflation is 2.5, and as $\beta \rightarrow 0$, it approaches 5, the Nash equilibrium of the model. This happens because as β decreases the policymaker cares less about long-run results and is more concerned about current outcomes. In the limiting case of $\beta = 0$ he cares only about today, so the discretionary result is obtained. Also, when β is 1 but b_1 is different from 1, the rule prescribes an inflation level different from zero: with $b_1 < 1$, inflation is set at a positive level, since people do not fully adapt their expectations to the observed values: the induction hypothesis is not satisfied, and the policymaker by generating a constant level of inflation can always "surprise" people and obtain a favorable result in terms of unemployment.¹⁶ If instead $b_1 > 1$, the same argument leads the

¹⁵This result requires the initial condition $\pi_0 = 0$. Otherwise, with initial conditions that imply the economy was on a higher inflation path, inflation will converge to 0 asymptotically.

¹⁶The levels of inflation that our policy prescribes for different values of b_1 , calculated with $a_0 = 5$ (the natural rate of unemployment), $a_1 = a_2 = 1$ and $\beta = 1$ are: inflation = 0 with $b_1 = 1$; inflation = 0.4950 with $b_1 = 0.9$; inflation = 2 with $b_1 = 0.5$. As $b_1 \rightarrow 0$, inflation $\rightarrow 2.5$, the optimal value when public's expectations about inflation are always at zero.

policymaker to set inflation at a negative level: since now people more than adapt their expectations to observed inflation (and deflation), any level of inflation would have a negative impact on unemployment, while by generating a constant level of deflation the authority can always "surprise" people with a level of deflation milder than expected, which produces a positive effect on unemployment. Note, however, that in our economy people have in fact rational expectations: if the policymaker recognized also this aspect, the optimal policy rule would tell him to set inflation always at zero.

The story is different when the policymaker does not know the parameter values and has to learn them from data, by recurrently estimating his model: in this case the estimated parameter values change over time, changing the optimal level of inflation prescribed by the rule. When the policymaker uses a decreasing gain algorithm to update his beliefs, the self-confirming equilibrium is again the Nash equilibrium ($U = 5$, $\pi = 5$), with parameter values in equations (30) and (32) in the Appendix as follows: $a_0 = 10$, $a_1 = 1$, $a_2 = 0$ and $b_1 = 1$. When, instead, a constant gain algorithm is implemented, the system doesn't settle down to a fixed point equilibrium, and escape dynamics appear.

4 Escape dynamics and policy specification

In light of what we have discussed in the previous section, we try now to understand better under which conditions the escape dynamics emerge. In particular, we want to check if their occurrence depends on the model used by the policymaker in gaining information from the data and setting its policy.

Our economy is still represented by equation (1): unemployment departs from its natural level only when unexpected inflation is generated. Since expectations are rational, no exploitable trade-off is present in the economy: the private sector can predict, up to a random error, the level of inflation set by the policymaker, and unemployment fluctuates randomly around its natural level. Any policy rule will be correctly anticipated by the public, and will have no effects on real variables. If the policymaker knew the true structure of the economy, two situations would possibly arise: the Nash equilibrium, in case of discretionary policy; or the Ramsey outcome, if credible commitment or reputational considerations were allowed in the problem. But in our model the policymaker doesn't know the true structure of the economy and tries to infer it from data: this introduces new dynamics in the system. How these dynamics depend on the model specification is the

topic addressed here.

We start endowing the policymaker with a correctly specified model, one that takes duly into account the role of expectations. The true Phillips curve is still of the form (1), but now the policymaker recursively estimates an equation of the type

$$U_t = a_{0,t-1} - a_{1,t-1}\pi_t + a_{2,t-1}\pi_t^e + \epsilon_{1,t}^{17} \quad (13)$$

where expectations are derived adaptively according to

$$\pi_t^e = b_{1,t-1}\pi_{t-1} + b_{2,t-1}\pi_{t-2} + b_{3,t-1}\pi_{t-3} \quad (14)$$

with the parameters recursively estimated from data, regressing inflation on lagged components.

The policymaker uses (13) as constraint in the minimization problem (5),¹⁸ and taking expectations as given obtains the optimal policy

$$\hat{\pi}_t = \frac{a_{0,t-1}a_{1,t-1}}{1 + a_{1,t-1}^2} + \frac{a_{1,t-1}a_{2,t-1}}{1 + a_{1,t-1}^2}\pi_t^e \quad (15)$$

The resulting inflation is shown in Figure 3, together with the values of the parameter $a_{1,t-1}$ and of the sum $-a_{1,t-1} + a_{2,t-1}$ in (13).¹⁹ The evolution of inflation follows exactly the evolution of the value of a_1 . Since the policymaker takes expectations as given, what matters is not the sum $-a_{1,t-1} + a_{2,t-1}$, representing the long-run trade-off in this adaptive model, but just $a_{1,t-1}$, which gives the short-run trade-off. When the Central Bank learns a model in which expectations are explicitly taken into account, and applies a discretionary policy, the sharp escape dynamics that characterize the vindication story do not arise. The policymaker sets inflation always at a positive value, well above the Ramsey zero-inflation level, with the specific magnitude at each time determined by the perceived short-run trade-off as given by the estimated $a_{1,t-1}$. Only if a_1 went to zero the policymaker would perceive that no trade-off is present in the data and set inflation equal to zero, but a_1 never goes to zero (see Figure 3). Thus, a policymaker following the advice coming from econometric policy evaluation practice with this type of policy rule would not have generated the sharp reduction in inflation observed during the 1980s.

¹⁷Equation (13) is a generalization of the expectations-augmented Phillips curve, one that allows the coefficients on inflation and expected inflation to differ from each other.

¹⁸For the derivation of the policy rule, see the Appendix, Case 2.

¹⁹The initial values for the simulations are set as follow: $a_{0,0} = 3$, $a_{1,0} = 1$, $a_{2,0} = 1$, $b_{1,0} = 1/3$, $b_{2,0} = 1/3$, $b_{3,0} = 1/3$. The other parameters are set as in the basic case.

We consider now a different situation: as before, the policymaker estimates a correctly specified Phillips curve, but now, in designing its policy, the authority considers the effects that current inflation will have on future expectations.²⁰ The optimal level of inflation to be set today depends therefore on future optimal policy, as discussed above. Once this intertemporal links are taken into account, the escape dynamics reappear. Figure 4 shows the policy implemented²¹ and the value for the sum $-a_1 + a_2b_1$.²² This sum comes from substituting the equation for expectations into the Phillips curve when deriving the policy. When this value is zero, no long run trade-off is perceived and the inflation is set to zero. A short-run trade-off is always present, since a_1 is systematically different from 0, but the policymaker does not try to exploit it. When the effects of policy on expectations are considered, long-run considerations, summarized by the induction hypothesis, guide the policymaker decisions. When this hypothesis is satisfied, no long run trade-off is perceived and the policymaker drives inflation to zero.

Comparing these two cases, it seems that what really matters is not the form of the model used for estimations but the way in which policy is derived. Even when the policymaker estimates a correctly specified model, he uses an adaptive proxy for inflation expectations: in this way he always finds a short run trade-off, since it takes time for expectations to adapt. Thus, if he cares only about today's outcome, the policymaker tries to exploit such a perceived trade-off and generates persistently a positive level of inflation. If instead in deriving his policy the CB takes into account long-run considerations, then what matters is the existence of a long run trade-off. When the estimated coefficients satisfy the induction hypothesis, no long-run trade-off is perceived, and inflation is driven to zero. The escape dynamics thus appear. The economy tends to go to its only self-confirming equilibrium, the Nash one, but this is not a stable outcome: stochastic variability in the data allow the policymaker to learn that no long run trade-off is really present and lead him to drive the inflation to zero.

This suggests that a policymaker that estimates a reduced form model, but cares only about today's outcome, would not generate escape dynamics.²³ And indeed our simulations reveal that this is the case. As Figure 5

²⁰For the derivation of the policy, see the Appendix, Case 3. An alternative procedure to find the optimal policy is by solving a dynamic programming problem, as is done by Sargent (1999). We tried also this way, and obtained similar results.

²¹Actual inflation is given by this policy plus a zero-mean random component.

²²The initial values for the simulations are set as follow: $a_{0,0} = 3$, $a_{1,0} = 1$, $a_{2,0} = 1$, $b_{1,0} = 1$. Other parameters as before, with the addition of β , now relevant, set at 0.98.

²³See Appendix, Case 4, for the specification of this case.

shows, the inflation path matches closely that of the parameter on current inflation in the estimated Phillips curve.²⁴

5 Discussion

Given the examples presented, we now try to provide some intuitions for the different results obtained. In the basic case presented at first, the model is static and short-run is indistinguishable from long-run, since no dynamics enter into the model. The policymaker sets inflation equal to zero when the coefficient on current inflation in the perceived Phillips curve goes to zero. Since in this mode it is not possible to distinguish long-run from short-run considerations, a parameter value of zero on current inflation means at the same time that neither long nor short run trade-off exist. The induction hypothesis, usually referring to a long-run property of the model (requiring a stable policy over many periods) is here given by the single parameter on current inflation, thus referring to a short run property.

Using dynamic models, the distinction between short run and long run becomes meaningful. When the policymaker implements short-run policies, taking expectations as given, what is needed in order for him to set inflation to zero is a coefficient on current inflation (the short-run trade-off) equal to zero. But data never provide such a value. When instead the policymaker considers long run effects, what is needed to induce him to drive inflation to zero is the fulfillment of the induction hypothesis. At times, data show this property, leading to disinflationary episodes.

When in Sargent's work the authority estimates a reduced form Phillips curve with lagged inflation and uses it to find the optimal policy, the policymaker implicitly takes into account the long run effects of the policy, through the lagged components in future Phillips curves that enter as constraints in the problem. The authority is in effect considering the effects of his policy on expectations, just not explicitly modeling them.²⁵

The two things must not be confused: reduced form vs. structural form models and short run vs. long run considerations. As remarked famously by Lucas (1976), but already noted before by members of the Cowles commission in the 1950s, in order to carry out policy analysis, it is necessary to have a structural model, one invariant to policy interventions. In Sar-

²⁴Initial values for the simulations are set as follow: $d_{1,0} = 3$, $d_{2,0} = 1$, $d_{3,0} = 0$, $d_{4,0} = 0$, $d_{5,0} = 0$. All the rest as before.

²⁵Sargent (1999) doesn't provide an analytical solution for the policy rule in the dynamic model, so we can not see explicitly how that policy is derived.

gent's work, the policymaker uses a reduced form model, which can not be regarded as policy invariant, but reestimates it over time, allowing for parameter drifting. In some sense, this procedure fixes the flaw of using a reduced form model.²⁶ This explains why, in this setting, using a reduced form or a structural form does not seem to make any difference.

What matters is the type of policy implemented: if the policymaker is concerned only by short-run effects, high inflation is bound to prevail, no matter what kind of model he uses for his policy analysis. Even knowing the true model, inflation would be at the non-optimal Nash equilibrium level. If instead long-run considerations enter the picture, the policymaker is led to reduce inflation: periodically, if an estimated Phillips curve is utilized (either in reduced or structural form), or more consistently, if theoretical considerations about the NRH are the guide for his actions.

In this setting, as said, using a reduced form or structural form model does not make much difference in terms of policy recommendations, because the policymaker continuously reestimates his model, in this way accounting for parameter drifting. Structural forms are supposed to be policy invariant, thus suitable for policy analysis (though not necessarily for long-run forecasts, since other sources of parameter drifting may be present); reduced form models instead are not meant to be policy invariant, their parameters varying with variation in policy (at least as long as rational expectations are involved).²⁷ Continuously reestimating a model compensates, to some degree, the problem of policy-induced parameter drifting in the Phillips curve. This is related to what Lucas (1976) suggested as a possible explanation for the success of short-term econometric forecasting: practising econometricians kept revising their estimated relations used for econometric forecasting, thus practically accounting for non policy-induced parameter drifting.

²⁶Note, anyway, that this procedure doesn't fix the flaw coming from the ignorance of the policymaker. Not knowing the true parameters, he is induced by the estimated values to set inflation at an higher level than the one he would obtain if, knowing the true values, the CB followed our rule 2 with long run considerations (here assuming, quite sensibly, that expectations fully adapt to actual outcomes).

²⁷Evans and Ramey (2006) show that the Lucas Critique applies to a broader set of models than the one characterized by rational expectations: in particular, models with adaptive expectations, provided that the updating parameter is optimally tuned, are also subject to the Critique.

6 Conclusions

In this work we have reconsidered the vindication story proposed by Sargent (1999) as a possible explanation for the burst in U.S. inflation during the late 1960s and 70s and for its subsequent conquest, trying to qualify its relevance. We have shown that the escape dynamics that lie at the core of that story do not depend on the econometric specification of the model the policymaker uses to gain information from data, but rather depend on the time-horizon of the policy, i.e., on whether the authority is concerned only with short-run or also with long-run results. There is no relation between the two aspects: long-run considerations, as well as short-run ones, can be carried out either in structural or in reduced form model, as long as those models are recurrently estimated to account for parameter drifting.

When the policymaker estimates a distributed lag Phillips curve, past inflation is explicitly present in the equation, and thus it is natural for him to take into account the effects of current policy on future outcomes. When instead the authority estimates a structural form, he needs to find a proxy for the expectational terms that explicitly enter into his estimated model. In our examples, the CB models private sector's expectations in an adaptive way, thus recognizing the impact of past policies on them. In addition, the CB recurrently estimates the adaptive form of those expectations, thus accounting for parameter shifting that changes in policy are likely to produce, even if expectations are not fully rational.

First identifying expectations and then, given those expectations, estimating a Phillips curve, or directly estimating a reduced form Phillips curve does not make any difference in terms of economic outcomes. What does make a difference is how the policymaker uses his estimated equations to derive his policy, and in particular whether he considers the short-run trade-off or the long-run one.

The vindication story of Sargent (1999) requires a policymaker that cares about long-run considerations and that is able to resist the temptation of a short-sighted management of the economy. In this sense, it is the same requirement necessary for the alternative story, the triumph of the NRH, to hold. Also in that case the policymaker, once understood that no permanent trade off is available, must be able to refrain himself from generating surprise in inflation in order to temporarily reduce unemployment.²⁸

As long as a policymaker is prone to exploit the short-run trade-off,

²⁸For a discussion of the relevance of the inflationary bias, and why this may not prevent a Central Bank from implementing the time-inconsistent optimal policy, see McCallum (1997).

there is no hope to obtain a low-inflation outcome. When he understands the long-run implications of his policy, better outcomes can be achieved: if this is done through econometric practice, the escape phenomenon is what we expect to see; if instead this understanding comes from more sound economic theory, the Ramsey, zero inflation result is possible.

7 Appendix

7.1 Derivation of policy: Case 1

The problem the policymaker has to solve in order to find the optimal policy is the following:

$$\min_{\hat{\pi}} E_{t-1} \sum_{t=1}^{\infty} \beta^{t-1} (\pi_t^2 + U_t^2) \quad (16)$$

$$s.t. U_t = a_{0,t-1} - a_{1,t-1} \pi_t + \epsilon_{1,t} \quad (17)$$

$$\pi_t = \hat{\pi}_t + \epsilon_{2,t} \quad (18)$$

Here, as well as in Case 2, Case 3 and Case 4, parameters in (17), once estimated, are taken as known for sure, as explained before in the paper. Substituting (17) and (18) in (16) and noting that no intertemporal dependence is involved, the problem simplifies:

$$\min_{\hat{\pi}} E_{t-1} [(\hat{\pi}_t + \epsilon_{2,t})^2 + (a_{0,t-1} - a_{1,t-1}(\hat{\pi}_t + \epsilon_{2,t}) + \epsilon_{1,t})^2] \quad (19)$$

whose FOC is

$$E_{t-1} [2(\hat{\pi}_t + \epsilon_{2,t}) + 2(a_{0,t-1} - a_{1,t-1}(\hat{\pi}_t + \epsilon_{2,t}) + \epsilon_{1,t})(-a_{1,t-1})] = 0, \quad (20)$$

which gives the optimal policy

$$\hat{\pi}_t = \frac{a_{0,t-1} a_{1,t-1}}{1 + a_{1,t-1}^2}. \quad (21)$$

7.2 Derivation of policy: Case 2

The policy problem is given by

$$\min_{\hat{\pi}} E_{t-1} \sum_{t=1}^{\infty} \beta^{t-1} (\pi_t^2 + U_t^2) \quad (22)$$

$$s.t. U_t = a_0 - a_1\pi_t + a_2\pi_t^e + \epsilon_{1,t} \quad (23)$$

$$\pi_t = \hat{\pi}_t + u_t \quad (24)$$

Here the policymaker takes public expectations as given, and does not consider the effect of his own policy on them. Substituting (23) and (24) in (22) we get the problem

$$\min_{\hat{\pi}_t} E_{t-1} \sum_{t=1}^{\infty} \beta^{t-1} ((\hat{\pi}_t + u_t)^2 + (a_0 - a_1(\hat{\pi}_t + u_t) + a_2\pi_t^e + \epsilon_{1,t})^2). \quad (25)$$

The FOC is

$$E_{t-1}[2(\hat{\pi}_t + u_t) + 2(a_0 - a_1(\hat{\pi}_t + u_t) + a_2\pi_t^e + \epsilon_{1,t})(-a_1)] = 0 \quad (26)$$

which implies

$$\hat{\pi}_t - a_0a_1 + a_1^2\hat{\pi}_t - a_1a_2\pi_t^e = 0. \quad (27)$$

The optimal policy therefore is

$$\hat{\pi}_t = \frac{a_0a_1}{1 + a_1^2} + \frac{a_1a_2}{1 + a_1^2}\pi_t^e \quad (28)$$

where $\pi_t^e = b_1\pi_{t-1}$.

7.3 Derivation of policy: Case 3

For simplicity of exposition, here and in the following Cases we drop the time subscript from the parameters. Parameters must be understood to be the most recent estimates available to the policymaker. For this reason, when the problem to be solved is dynamic, they are the same in current and future constraints, since the policymaker when solving his policy problem at any time, takes the more recent estimates to hold forever in the future.

The problem to be solved by the policymaker is as follow:

$$\min_{\hat{\pi}} E_{t-1} \sum_{t=1}^{\infty} \beta^{t-1} (\pi_t^2 + U_t^2) \quad (29)$$

$$s.t. U_t = a_0 - a_1\pi_t + a_2\pi_t^e + \epsilon_{1,t} \quad (30)$$

$$\pi_t = \hat{\pi}_t + u_t \quad (31)$$

$$\pi_t^e = b_1\pi_{t-1} \quad (32)$$

Substituting (32) in (30) and the resulting expression, together with (31), in (29) we get

$$\min_{\hat{\pi}_t} E_{t-1} \sum_{t=1}^{\infty} \beta^{t-1} ((\hat{\pi}_t + u_t)^2 + (a_0 - a_1(\hat{\pi}_t + u_t) + a_2 b_1 \pi_{t-1} + \epsilon_{1,t})^2) \quad (33)$$

The FOC is

$$E_{t-1} [2(\hat{\pi}_t + u_t) + 2(a_0 - a_1(\hat{\pi}_t + u_t) + a_2 b_1 \pi_{t-1} + \epsilon_{1,t})(-a_1) + 2\beta(a_0 - a_1(\hat{\pi}_{t+1} + u_{t+1}) + a_2 b_1(\hat{\pi}_t + u_t) + \epsilon_{1,t+1})(a_2 b_1)] = 0, \quad (34)$$

which can be simplified in

$$E_{t-1} [\hat{\pi}_t + (a_0 - a_1 \hat{\pi}_t + a_2 b_1 \pi_{t-1})(-a_1) + \beta(a_0 - a_1 \hat{\pi}_{t+1} + a_2 b_1 \hat{\pi}_t)(a_2 b_1)] = 0 \quad (35)$$

or

$$E_{t-1} [\hat{\pi}_t - a_0 a_1 + a_1^2 \hat{\pi}_t - a_1 a_2 b_1 \pi_{t-1} + \beta a_0 a_2 b_1 - \beta a_1 a_2 b_1 \hat{\pi}_{t+1} + \beta a_2^2 b_1^2 \hat{\pi}_t] = 0, \quad (36)$$

and solving for the policy instrument we obtain:

$$\hat{\pi}_t = \frac{1}{1 + a_1^2 + \beta a_2^2 b_1^2} [(a_0 a_1 - \beta a_0 a_2 b_1) + (a_1 a_2 b_1) \pi_{t-1} + \beta a_1 a_2 b_1 E_{t-1} \hat{\pi}_{t+1}] \quad (37)$$

For simplicity of exposition we write

$$k_0 = \frac{a_0 a_1 - \beta a_0 a_2 b_1}{1 + a_1^2 + \beta a_2^2 b_1^2} \quad (38)$$

$$k_1 = \frac{a_1 a_2 b_1}{1 + a_1^2 + \beta a_2^2 b_1^2} \quad (39)$$

$$k_2 = \frac{\beta a_1 a_2 b_1}{1 + a_1^2 + \beta a_2^2 b_1^2} \quad (40)$$

and the policy rule becomes:

$$\hat{\pi}_t = k_0 + k_1 \pi_{t-1} + k_2 E_{t-1} \hat{\pi}_{t+1}. \quad (41)$$

Equation (41) says that the optimal policy today depends on the optimal policy the authority expects to implement in the next period.

The true Phillips curve

$$U_t = U^* + \theta(\pi_t - E_{t-1} \pi_t) + \nu_{1,t} \quad (42)$$

and (41), together with condition (31), give solution for U_t and π_t . But (41) is clearly independent from (42) and gives solution for $\hat{\pi}_t$. To find an explicit form for $\hat{\pi}_t$ we look for the MSV solution, using the undetermined coefficient procedure.

Given that the only state variable in (41) is lag inflation, we guess a solution of the form

$$\hat{\pi}_t = \phi_0 + \phi_1 \pi_{t-1}. \quad (43)$$

Therefore, using condition (31), it follows that

$$E_{t-1} \hat{\pi}_{t+1} = \phi_0 + \phi_1 E_{t-1} \pi_t = \phi_0 + \phi_1 \hat{\pi}_t = \phi_0 + \phi_1 (\phi_0 + \phi_1 \pi_{t-1}), \quad (44)$$

and for our guess to be correct it must be that

$$\phi_0 + \phi_1 \pi_{t-1} = k_0 + k_1 \pi_{t-1} + k_2 (\phi_0 + \phi_1 (\phi_0 + \phi_1 \pi_{t-1})) \quad (45)$$

which implies

$$\phi_0 = k_0 + k_2 \phi_0 + k_2 \phi_1 \phi_0 \quad (46)$$

and

$$\phi_1 = k_1 + k_2 \phi_1^2. \quad (47)$$

From the second equation we have

$$\phi_1 = \frac{1 \pm \sqrt{1 - 4k_1 k_2}}{2k_2} \quad (48)$$

and from the first

$$\phi_0 = \frac{k_0}{1 - k_2(1 + \phi_1)}. \quad (49)$$

The MSV solution²⁹ is given by³⁰

$$\phi_1 = \frac{1 - \sqrt{1 - 4k_1 k_2}}{2k_2} \quad (50)$$

$$\phi_0 = \frac{k_0}{1 - k_2 \left(1 + \frac{1 - \sqrt{1 - 4k_1 k_2}}{2k_2}\right)}. \quad (51)$$

²⁹The MSV solution is the bubble-free solution, the one based only on fundamentals. See McCallum (1999).

³⁰The reason why this is the solution selected by the MSV criterion is that when $b_1 = 0$, and therefore $k_1 = 0$, it must be that $\phi_1 = 0$. The non MSV solution is instead given by $\phi_1 = \frac{1 + \sqrt{1 - 4k_1 k_2}}{2k_2}$ and leads to an explosive path for inflation.

7.4 Derivation of policy: Case 4

Since the policymaker in this case is concerned only about today's outcome, his problem is as follows:

$$\min_{\hat{\pi}_t} E_{t-1}(\pi_t^2 + U_t^2) \quad (52)$$

$$s.t. U_t = d_0 - d_1\pi_t + d_2\pi_{t-1} + d_3\pi_{t-2} + d_4\pi_{t-3} + \epsilon_{1,t} \quad (53)$$

$$\pi_t = \hat{\pi}_t + u_t. \quad (54)$$

Substituting (53) and (54) in (52) we obtain

$$\min_{\hat{\pi}_t} E_{t-1}((\hat{\pi}_t + u_t)^2 + (d_0 - d_1(\hat{\pi}_t + u_t) + d_2\pi_{t-1} + d_3\pi_{t-2} + d_4\pi_{t-3} + \epsilon_{1,t})^2). \quad (55)$$

The optimality condition is given by

$$E_{t-1}[2(\hat{\pi}_t + u_t) + 2(d_0 - d_1(\hat{\pi}_t + u_t) + d_2\pi_{t-1} + d_3\pi_{t-2} + d_4\pi_{t-3} + \epsilon_{1,t})(-d_1)] = 0, \quad (56)$$

which implies

$$\hat{\pi}_t - d_0d_1 + d_1^2\hat{\pi}_t - d_1d_2\pi_{t-1} - d_1d_3\pi_{t-2} - d_1d_4\pi_{t-3} = 0 \quad (57)$$

and therefore the optimal policy

$$\hat{\pi}_t = \frac{d_0d_1}{1 + d_1^2} + \frac{d_1d_2}{1 + d_1^2}\pi_{t-1} + \frac{d_1d_3}{1 + d_1^2}\pi_{t-2} + \frac{d_1d_4}{1 + d_1^2}\pi_{t-3}. \quad (58)$$




Figure 1. Case 1: Inflation and unemployment.




Figure 2. Case 1: Evolution of the estimated parameter values in equation (17).




Figure 3. Case 2: Inflation and estimated parameter values in equation (23).




Figure 4. Case 3: Policy and estimated parameter values in equation (30) (after substituting in equation 32).




Figure 5. Case 4: Inflation and estimated parameter values in equation (53).

References

- [1] Barro, Robert J. and David B. Gordon (1983a), "Rules, discretion and reputation in a model of monetary policy", *Journal of Monetary Economics* 12, 101-121.
- [2] Barro, Robert J. and David B. Gordon (1983b), "A positive theory of monetary policy in a natural rate model", *Journal of Political Economy* 91, 589-610.
- [3] Beck, Gunter W. and Volker Wieland (2002), "Learning and control in a changing economic environment", *Journal of Economic Dynamics and Control* 26, 1359-1377.
- [4] Cho, In-Koo, Noah Williams and Thomas J. Sargent (2002), "Escaping Nash inflation", *Review of Economic Studies* 69, 1-40.
- [5] Evans, George W. and Seppo Honkapohja (2001), *Learning and expectations in macroeconomics*, Princeton University Press.
- [6] Evans, George W. and Garey Ramey (2006), "Adaptive expectations, underparametrization and the Lucas critique", *Journal of Monetary Economics* 53, 249-264.
- [7] Friedman, Benjamin (1979), "Optimal expectations and the extreme information assumption of 'rational expectations' macromodels", *Journal of Monetary Economics* 5, 23-41.
- [8] Gerali, Andrea and Francesco Lippi (2002), "On the 'conquest' of inflation", Banca d'Italia, Temi di discussione del Servizio Studi, 444.
- [9] Ireland, Peter N. (1999), "Does the time-consistency problem explain the behavior of inflation in the United States?". *Journal of Monetary Economics* 44, 279-291.
- [10] Kendrick, David A. (2002), *Stochastic control for economic models*, Second edition, Version 2.00, on-line.
- [11] King, Robert G. and Mark W. Watson (1994), "The post-war U.S. Phillips curve: A revisionist econometric history", *Carnegie-Rochester Conference Series on Public Policy* 41, 157-219.
- [12] King, Robert G. and Mark W. Watson (1997), "Testing long-run neutrality", *Economic Quarterly* 83/3, Federal Reserve Bank of Richmond.

- [13] Kydland, Finn and Edward C. Prescott (1977), "Rules rather than discretion: the inconsistency of optimal plans", *Journal of Political Economy* 85, 473-491.
- [14] Ljungqvist, Lars and Thomas J. Sargent (2000), *Recursive Macroeconomic Theory*, MIT Press.
- [15] Lucas, Robert E. (1976), "Econometric policy evaluation: A critique", *Carnegie-Rochester Conference Series on Public Policy*, 19-46.
- [16] Marschak, Jacob (1953), "Economic measurements for policy and prediction", in William C. Hood and Tjalling G. Koopmans eds., *Studies in Econometric Methods*, Cowles Commission Monograph 14, New York: Wiley, 1-26.
- [17] Maybeck, Peter S. (1979), *Stochastic Models, Estimation and Control*, Academic Press, Vol.1, Ch. 1.
- [18] McCallum, Bennett T. (1994), "Identification of inflation-unemployment tradeoffs in the 1970s. A comment", *Carnegie-Rochester Conference Series on Public Policy* 41, 231-241.
- [19] McCallum, Bennett T. (1997), "Crucial issues concerning central bank independence", *Journal of Monetary Economics* 39, 99-112.
- [20] McCallum, Bennett T. (1999), "Role of the minimal state variable criterion in rational expectations models", *International Tax and Public Finance* 6, 621-639. Also in Isard, P., Assaf Razin and Andrew K. Rose (eds.) *International Finance and Financial Crises: Essays in Honor of Robert P. Flood Jr.*, Kluwer Academic Press.
- [21] Muth, John F. (1960), "Optimal properties of exponentially weighted forecasts", *Journal of the American Statistical Association* 55, 299-306.
- [22] Sargent, Thomas J. (1993), *Bounded Rationality in Macroeconomics*, Clarendon Press.
- [23] Sargent, Thomas J. (1999), *The Conquest of American Inflation*, Princeton University Press.
- [24] Sims, Christopher A. (1988), "Projecting policy effects with statistical models", Paper prepared for the Latin American Meetings of the Econometric Society, San Jose, Costa Rica, August 1988.

- [25] Taylor, John B. (1996), "How should monetary policy respond to shocks while maintaining long-run price stability? - Conceptual Issues", in Proceeding of the 1996 Symposium *Achieving Price Stability*, sponsored by the Federal Reserve Bank of Kansas City.
- [26] Tetlow, Robert J. and Peter von zur Muhelen (2004), "Avoiding Nash inflation", *Review of Economic Dynamics* 7, 869-899..
- [27] Welch, Greg and Gary Bishop (2002), "An introduction to the Kalman filter", Working paper.
- [28] Wieland, Volker (2000), "Learning by doing and the value of optimal experimentation", *Journal of Economic Dynamics and Control* 24, 501-534.