Imperfect Knowledge, Adaptive Learning and the Bias against Activist Monetary Policies

Alberto Locarno
Banca d'Italia

2006 MPRW in LA
Lima, November 7th

Outline of the presentation

- Objectives and findings
- Reference to the literature
- The theoretical model: features and properties
- Introducing adaptive learning: stability and speed of convergence
- Some simulation results
- Conclusions

Objectives and findings (1)

- **Two main issues:**

  How does the relaxation of the REH affect monetary policymaking?

  Does learning change the optimal degree of monetary policy discretion?

- **Basic ingredients:**

  Recurrent shifts in the structure of the economy make the REH untenable and call for some form of bounded rationality and learning.

  Incomplete knowledge precludes the possibility for the central bank to commit to a given policy, since the private sector cannot verify whether the policymaker is delivering on its promises.

  When the central bank has private information, an efficient way to reduce the welfare loss associated with discretion is to adopt policies that take the form of an inflation cap (AAK, 2005).

  Lexicographic preferences, which represent quite closely Buba's and ECB's mandates, promote policies that take the form of an inflation cap.

Objectives and findings (3)

Perfect knowledge
Rational Expectations → Commitment → Optimal policies

But:

Changing economic structure → Imperfect knowledge

Imperfect knowledge: Unverifiable actions, Learning → Commitment impossible and ineffective

→ Discretion → What policies work fine? How do they differ from those under commitment?

Private information creates an additional tension between time-consistency and discretion

Objectives and findings (4)

- **Main findings:**

  With incomplete knowledge and adaptive learning the incentives and constraints facing the monetary authority change and a bias towards conservatism arises.

  When agents and policymakers learn adaptively, additional inertia is introduced in the system, which makes costly for the central bank not to respond promptly and forcefully to (inflationary) shocks.

  A policy involving a cap on inflation is helpful in reducing output and inflation variability, but it is not uniformly the most efficient.

  The bias against stabilisation policies and the relative efficiency of alternative monetary strategies do not depend on the specific form of the learning process, namely on whether memory is finite (constant gain learning) or infinite (decreasing gain learning).

Reference to the literature (1)

- **The paper is related to the literature in several ways:**

it builds on the work by Orphanides and Williams (2002) to study how the economy responds to alternative monetary strategies when agents have bounded rationality and imperfect knowledge;

it models two-side learning in the vein of Evans-Honkapohja (2002), i.e. both the private sector and the monetary authority learn adaptively;

as in Terlizzese (1999) and Driffil-Rotondi (2003), it explicitly considers a a lexicographic preference ordering for the monetary authority;

it provides a quantitative assessment of the claim by Athey *et al.* (2003) that the optimal policy under discretion takes the form of an inflation cap.

Reference to the literature (2)

- The original contribution of the paper is to extend and generalise the findings by Orphanides and Williams. It does it in at least three different ways:
  1. it proves the claim that imperfect knowledge and adaptive learning induce a bias against stabilisation policies also in models in which not only inflation expectations but also output expectations matter;
  2. it provides evidence that, if one drops the assumption of rational expectations, it is of primary importance for the policymaker to avoid that target and expected inflation diverge also in models without intrinsic dynamics;
  3. it uses the theoretical insights on the optimal degree of monetary policy discretion to test whether society can increase welfare by appointing a policymaker whose preferences are lexicographic;
  4. it studies the impact of imperfect knowledge on stabilisation policies under a set of alternative learning mechanisms, without relying exclusively on constant-gain algorithms.

The theoretical model (1)

- The model has two basic ingredients: (1) unobservability of the supply shock; (2) unknown and time varying output-inflation trade-off.

1. The economy is characterised by a Lucas-type supply curve:

$$y = \alpha \left(\pi - \pi^{e}\right) + \varepsilon$$

   the output gap depends on inflation surprises and on an unobservable mean-zero supply shock ε, uniformly distributed on the interval $[-\mu, \mu]$.

2. Because of frictions, the desired level of output for the central bank is $k>0$.

3. π is the monetary policy instrument, controlled without errors.

4. A signal z, observed only by the policymaker, conveys noisy information on the output shock ε. By assumption, $z = \varepsilon + \xi$, with $\xi \sim U[-\mu, \mu]$.

5. The output-inflation trade-off α is random and changes every period. It is assumed independent of all other shocks and $IID\left(\overline{\alpha}, \sigma_{\alpha}^{2}\right)$.

The theoretical model (2)

5. Information in the economy flows as follows:

π^e_t	z_t	π_t	y_t
agents form inflation expectations	output signal observed	central bank chooses inflation	output realised

The signal *z* materialises before the central bank chooses the inflation rate but after private agents set their inflation expectations. The information advantage of the monetary authority creates a role for stabilisation policies.

In setting the inflation rate so as to (partially) offset the output shock, the central bank solves a signal extraction problem. Given the nature of the stochastic shocks, $E(\varepsilon|z)$, i.e. the expected value of ε conditional on the value of *z*, turns out to be equal to $z/2$. Notice that if the loss function is quadratic, the signal extraction problem and the optimal setting of the policy instrument the can be kept separated and solved sequentially.

The theoretical model (3)

6. Society's loss function is quadratic. The central bank is assumed to have either lexicographic or quadratic utility. In the principal-agent approach, society, whose preferences are quadratic, can appoint a central banker whose loss function is different from its own, so as to reduce the discretionary problem and maximise welfare.

   Though not true in general, in the simplified case of two objectives lexicographic preferences can be represented by a function. The central bank problems is:

$$\min_{\pi} \quad \frac{1}{2}E\left(y-k\right)^2$$
$$s.t. \quad \begin{cases} \pi \le \bar{\pi} \\ y = \alpha\left(\pi - \pi^2\right) + \varepsilon \end{cases}$$

   where ***k>0*** is the target level of output and $\bar{\pi}$ is the upper bound on inflation. To allow for a closed form solution, it is assumed that ***k= μ/6***.

   With time-separable preferences, the problem is static and involves no trade-off between current and future utility, so that the optimal policy does not depend on the strategic interactions between the central bank and the private sector.

The theoretical model (4)

The optimal policies for lex. preferences turns out to be:

$$\pi = \begin{cases} \pi^e - \dfrac{\overline{\alpha}}{\overline{\alpha}^2 + \sigma_\alpha^2}\left(\dfrac{z}{2} - k\right) \;=\; \pi^e - \dfrac{z}{\phi} + \dfrac{2k}{\phi} & \textit{if } z \geq 2k + \phi\left(\pi^e - \overline{\pi}\right) \\ \overline{\pi} & \text{otherwise} \end{cases}$$

where $\phi^{-1} \equiv \dfrac{1}{2}\dfrac{\overline{\alpha}}{\overline{\alpha}^2 + \sigma_\alpha^2}$ **. The equilibrium is non-cooperative Nash.**

Remarks:

the optimal policy takes the form of an inflation cap;

uncertainty on α ***and unobservability of*** ε ***attenuate the response of the policymaker;***

when ε ***is not too negative, the monetary authority can stabilise output more effectively than under quadratic preferences;***

$\pi^e = \overline{\pi} - \dfrac{2k}{\phi} < \overline{\pi}$ ***, i.e. the existence of an upper bound on inflation contributes to keep price dynamics moderate;***

the upper bound on inflation helps reducing the extent of the inflationary bias: the lower the upper bound, the lower the expected inflation;

a lower bound on inflation does not change the equilibrium properties of the model: the central bank has an incentive to inflate and not to disinflate.

The theoretical model (5)

If the monetary policymaker is instead endowed with quadratic preferences, with ***β*** being the weight of the inflation objective, optimal policy takes the following form:

$$\begin{aligned}\pi &= \frac{\bar{\alpha}^2+\sigma_\alpha^2}{\bar{\alpha}^2+\sigma_\alpha^2+\beta}\pi^e + \frac{\bar{\alpha}}{\bar{\alpha}^2+\sigma_\alpha^2+\beta}\left(k-\frac{z}{2}\right) = \rho\pi^e + \frac{\rho}{\phi}(2k-z) \\ &= \pi^e - \frac{\rho}{\phi}z\end{aligned}$$

where $\rho \equiv \frac{\bar{\alpha}^2+\sigma_\alpha^2}{\bar{\alpha}^2+\sigma_\alpha^2+\beta}$ and $\pi^e = \frac{\bar{\alpha}}{\beta}k$. Depending on the value of ***β***, such a policy can provide more or less output stabilisation compared with the one minimising a lexicographic preference ordering.

A policy minimising quadratic losses represents the obvious benchmark for assessing the properties of “inflation cap” policies.

Introducing adaptive learning (1)

- When the economy is subject to recurrent shifts, people have to figure out how the environment is changing: in such a context, the REH is not suitable.
- Adaptive learning represents a specific form of bounded rationality. It provides an asymptotic justification for the REE and allows to neglect non-learnable solutions in models with multiple equilibria.
- The central idea is that the each period *t* agents have a PLM they use to make forecasts. The latter influence decisions at time *t* and yield the temporary equilibrium (the ALM). The ALM provides a new data point that agents use to update forecasts for period *t+1*. The learning dynamics continues with the same steps in every period and represents a driving force of the system.
- Projection coefficients update originates a system of recursive equations, describing the mapping between the PLM and the ALM. Convergence of the process is studied by means of the associated ODE: stability holds when the (real part of the) eigenvalues of the Jacobian are negative (E-stability).
- On the speed of convergence, the only analytical result is that root-*t* convergence applies when the largest eigenvalue of the ODE is less the *-½*.

Introducing adaptive learning (2)

If we assume that agents learn adaptively, two cases arise: (1) the central bank has RE while the private sector has imperfect knowledge; (2) both players learn. In the first case:

$$y = \alpha\left(\pi - \hat{E}^{P}\pi\right) + \varepsilon$$

$$\pi = \begin{cases} \min\left[\hat{E}^{CB}\pi - \dfrac{z}{\phi} + \dfrac{2k}{\phi}, \overline{\pi}\right] & \textit{(lex. pre ferences)} \\ \rho\hat{E}^{CB}\pi - \dfrac{\rho}{\phi}z + \dfrac{\rho}{\phi}2k & \textit{(quad. pre ferences)} \end{cases}$$

$\hat{E}^{P}\pi$ **represents the current estimate of the inflation rate of the private sector; the monetary authority, moving second and having rational expectations, sets** $\hat{E}^{CB}\pi = \hat{E}^{P}\pi$ **.**

At each period *t*, private agents have a PLM for inflation taking the form:

$$\hat{E}^{P}\pi = a_{P,t} = a_{P,t-1} + \frac{1}{t}\left(\pi_{t-1} - a_{P,t-1}\right)$$

The mapping between the PLM and the ALM generates a stochastic recursive algorithm, which is approximated by the following ODE:

$$\frac{d}{d\tau}a_P = h\left(a_P\right) = \lim_{t\to\infty} E\left(\pi_{t-1} - a_{P,t-1}\right) = \begin{cases} \dfrac{1}{2}\left(\overline{\pi} - a_P\right) - \dfrac{\mu}{6\phi} & \textit{(lex. preferences)} \\ \left(\rho - 1\right)a_P + \dfrac{\rho}{\phi}2k & \textit{(quad. preferences)} \end{cases}$$

Introducing adaptive learning (3)

- **For both policies, the fixed point of the ODE, namely $a_P = h^{-1}(0)$, coincides with the unique REE; the ODE is stable and hence asymptotically adaptive learning converges to the REE. The only condition required for stability is that the inflation objective enters the loss function of the policymaker.**

- **For reasonable parameterisation of the model, root-*t* convergence does not hold. What is required for increasing the speed of learning is that the monetary authority is highly inflation-averse and accordingly inflation is not too volatile.**

- **As in Orphanides-Williams (2003), the ability of private agents to forecast inflation depends on the monetary policy in place: more aggressive policies reduce the persistence of inflation, facilitates the formation of expectations, enhances economic stability and mitigates the influence of imperfect knowledge on the economy. This findings however is likely to be biased by the structure of the model, where most variability is due to fluctuations in inflation expectations.**

Introducing adaptive learning (4)

If the monetary authority doesn't know the moments of the output-inflation trade-off, it needs somehow to estimate them, to be able to set π properly. To account for central bank learning, the previous model must be augmented with a new set of recursive equations, which are the same irrespective of the policymaker's preferences. The set of recursive least squares equation is:

$$
\begin{aligned}
a_{P,t} &= a_{P,t-1} + \frac{1}{t}\left(\pi_{t-1} - a_{P,t-1}\right) \\
\hat{\alpha}_t &= \hat{\alpha}_{t-1} + R_{\pi,t-1}^{-1}\left(\pi_{t-1} - a_{P,t-1}\right)\frac{1}{t}\left[\left(y_{t-1} - \frac{z_{t-1}}{2}\right) - \hat{\alpha}_{t-1}\left(\pi_{t-1} - a_{P,t-1}\right)\right] \\
R_{y,t} &= R_{y,t-1} + \frac{1}{t}\left[\left(y_{t-1} - \frac{z_{t-1}}{2}\right)^2 - R_{y,t-1}\right] \\
R_{\pi,t} &= R_{\pi,t-1} + \frac{1}{t}\left[\left(\pi_{t-1} - a_{P,t-1}\right)^2 - R_{\pi,t-1}\right]
\end{aligned}
$$

The first equations, as before, captures private sector learning, while the others refer to the central bank's inference problem.

Introducing adaptive learning (5)

- **The central bank computes the statistics $R_{y,t}$ and $R_{\pi,t}$ as an intermediate step to estimate the optimal response ϕ^{1} (and $\rho\,\phi^{1}$) to the signal z.**

- **Since the central bank observes z, it knows $y - \frac{z}{2}$, the policy driven component of output, and can efficiently estimate $\bar{\alpha}$ by regressing it on the inflation surprise. Efficiency stems form the fact that $\varepsilon - \frac{z}{2}$ is orthogonal to the signal z.**

- **A biased estimate of $E(\alpha^2)$ can be obtained from the sample average of the squared policy-driven component of the output gap, scaled by the second moment of the inflation surprise:**

$$\frac{E\left(y - \frac{z}{2}\right)^2}{E\left(\pi - a_P\right)^2} = \frac{E\left(\alpha^2\right)E\left(\pi - a_P\right)^2 + E\left(\varepsilon - \frac{z}{2}\right)^2}{E\left(\pi - a_P\right)^2} = \bar{\alpha}^2 + \sigma_{\alpha}^2 + \frac{2\frac{\mu^2}{3}}{E\left(\pi - a_P\right)^2}$$

To eliminate the bias, the third term in the last equality is therefore subtracted from the ratio of $R_{y,t}$ to $R_{\pi,t}$.

Introducing adaptive learning (6)

- Regardless policymaker's preferences, the ODE of the system is recursive. $R_{\pi} \rightarrow E(\pi - a_P)^2$ from any starting point and the same happens for R_y. Provided R_{π} is invertible along the convergence path, $R_{\pi}^{-1}E(\pi - a_P)^2 \rightarrow 1$ and the differential equation for $\hat{\bar{\alpha}}$ may be assessed independently of the remaining part of the system.

- Conditional on $\hat{\bar{\alpha}}$, R_{π} and R_{π} approaching the true values, convergence to the REE of the private sector expectations is determined on the basis of the eigenvalues of the ODE for a_P only.

- The effect of preferences on the speed of convergence is not clear. If also the policymaker has imperfect knowledge, an additional layer of interactions between policy actions and economic outcomes arises and both the asymptotic and finite sample behaviour becomes relevant. In addition, when the learning process is disturbed by several shocks, only for large values of t the ODE becomes an acceptable approximation of the SRA and the asymptotic distribution is not of much help in understanding the properties of the system. In the case of multiple equilibria, the problem is even more serious, since with few *dof* large shocks can displace the RLS estimates outside the domain of attraction of the ODE.

Some simulation results (1)

- Model simulations are used to illustrate how learning affects the dynamics of the system.
- First, the performance of the forecasting rule is assessed by focusing on the mean and median of the inflation forecasts.
- Then the speed of convergence is studied according to the numerical procedure proposed by Marcet and Sargent.
- Finally, the output-inflation variability trade-off is computed and the role of policy regimes and learning is appraised.
- To test whether indeed a bias against stabilisation policies exists, additional simulations mimicking the impact on the economy of a string of negative shocks are run.
- As a final check on how much the results depend on the specific leaning mechanism, the case of perpetual leaning is also considered.

Some simulation results (2)

- Each experiment is based on 500 replications and covers a 2000 period interval.
- Subsamples of 500 periods are considered to estimate the convergence speed.
- Initial conditions for the lagged variables in the RLS algorithms are randomly chosen from the REE distribution.
- The initial 150 observations are excluded so as to reduce the impact of initial conditions.
- Under lexicographic preferences, it is assumed that the upper bound on inflation is chosen so as to drive π^e to zero.

Some simulation results (3)

- Table 2a shows (for a d.g.s.) a few summary statistics on agents' inflation forecasting model.
- For most value of β, a policymaker with lex. preferences (strategy 1) is more effective in anchoring expectations to the REE. The higher precision of the forecast is confirmed by the smaller standard deviation and the unbiasedness of a_p.
- For both strategies, the mean and the median concide.
- No strategy is better in terms of speeding the convergence process. Convergence is however nearly immediate and the estimate of δ are likely to be not entirely reliable.
- The welfare reduction of departing from the perfect knowledge benchmark is negligible.
- Policy prescription: when the economic environment is not too complex, deviations form the REE are not too costly and a benevolent government is better of by appointing a policymaker with lex. Preferences.

Some simulation results (4)

Table 2a - Single-agent learning

(unconstrained estimator - decreasing gain sequence)

	Lexicographic preferences		Quadratic preferences	
	RE	*T=2000*	*RE*	*T=2000*
β = 0.176				
mean a_P	0,0000	0,0000	0,0289	0,0289
median a_P	0,0000	0,0000	0,0289	0,0289
SD a_P	-	0,0002	-	0,0289
output variability	0,0084	0,0084	0,0074	0,0074
inflation variability	0,0022	0,0022	0,0291	0,0291
convergence speed	δ = 0.26874		δ = 0.32196	
β =1.0				
mean a_P	0,0000	0,0000	0,0051	0,0044
median a_P	0,0000	0,0000	0,0051	0,0045
SD a_P	-	0,0002	-	0,0044
output variability	0,0084	0,0084	0,0076	0,0076
inflation variability	0,0022	0,0022	0,0059	0,0054
convergence speed	δ = 0.26874		δ = 0.19389	
β =5.667				
mean a_P	0,0000	0,0000	0,0009	0,0009
median a_P	0,0000	0,0000	0,0009	0,0009
SD a_P	-	0,0002	-	0,0009
output variability	0,0084	0,0084	0,0086	0,0086
inflation variability	0,0022	0,0022	0,0017	0,0017
convergence speed	δ = 0.26874		δ = 0.31993	

no skewness

Low speed.
Reliable?

lex. preferences
work

welfare reduction
negligible

Some simulation results (5)

- Tables 3a and 4a show the results for the case when both players learn; the former deals with the "plain" RLS rule (UE), while the latter shows results for the case of constrained estimation (CE).
- The tables report results for the mean of a_p, the volatility of π and y and the estimate of the policy rule coefficient and convergence speed.
- A striking finding is that the loss in welfare caused by imperfect knowledge is modest, below 10%. This is remarkable, since the inference problem of the policymaker is quite convoluted. Possible role of the rejection rate.
- Strategy 1 is not uniformly superior, though preferable for low values of β. The speed of convergence is nearly the same in both cases and $\delta \approx ½$. For any value of β, the estimate of ρ/ϕ is surprisingly accurate.
- Under lex. preferences, estimates of ϕ are upward biased, though not in the CE case, which is responsible for some undesired fluctuation in output. The excessive inflation volatility is not due to the surprise component but to a_p.
- The bounds imposed on RLS work a a projection facility: the rejection rate in the CE case is substantially lower than in the UE case. In the latter, the complexity of the filtering problem often displaces the RLS algorithm outside the domain of attraction of the ODE. If the estimate of ϕ becomes large, the policymaker has no incentive to respond to the signal and changes in y mostly reflect the output shock: the data become uninformative and the RLS estimates get trapped far away from the true value.

Some simulation results (6)

Table 3a - Two-agent learning

(unconstrained estimator - decreasing gain sequence)

	Lexicographic preferences		Quadratic preferences	
	RE	*T=2000*	*RE*	*T=2000*
β = 0.176				
mean a_P	0,0000	-0,0005	0,0289	0,0283
policy rule coefficient	0,2642	0,1807	0,2508	0,2308
output variability	0,0084	0,0088	0,0074	0,0076
inflation variability	0,0022	0,0023	0,0291	0,0298
convergence speed	a_P: δ = 0.3351		a_P: δ = 0.3570	
	α: δ = 0.4951		α: δ = 0.2435	
	ψ: δ = 0.4609		ψ: δ = 0.4884	
β =1.0				
mean a_P	0,0000	-0,0005	0,0051	0,0048
policy rule coefficient	0,2642	0,1807	0,2029	0,1844
output variability	0,0084	0,0088	0,0076	0,0078
inflation variability	0,0022	0,0023	0,0059	0,0059
convergence speed	a_P: δ = 0.3351		a_P: δ = 0.2185	
	α: δ = 0.4951		α: δ = 0.2728	
	ψ: δ = 0.4609		ψ: δ = 0.4892	
β =5.667				
mean a_P	0,0000	-0,0005	0,0009	0,0008
policy rule coefficient	0,2642	0,1807	0,0975	0,0775
output variability	0,0084	0,0088	0,0086	0,0088
inflation variability	0,0022	0,0023	0,0017	0,0016
convergence speed	a_P: δ = 0.3351		a_P: δ = 0.4080	
	α: δ = 0.4951		α: δ = 0.3642	
	ψ: δ = 0.4609		ψ: δ = 0.4891	

Some simulation results (7)

upward bias
for lex.pref.

lex. pref. not
uniformly best

bounds work as a
projection facility

Table 4a - Two-agent learning

(constrained estimator - decreasing gain sequence)

	Lexicographic preferences		Quadratic preferences	
	RE	*T=2000*	*RE*	*T=2000*
β = 0.176				
mean a_P	0,0000	-0,0003	0,0289	0,0267
policy rule coefficient	0,2642	0,2365	0,2508	0,2411
output variability	0,0084	0,0087	0,0074	0,0075
inflation variability	0,0022	0,0025	0,0291	0,0271
convergence speed	a_P: δ = 0.4193		a_P: δ = 0.2972	
	α: δ = 0.4670		α: δ = 0.4557	
	ψ: δ = 0.4765		ψ: δ = 0.4892	
β =1.0				
mean a_P	0,0000	-0,0003	0,0051	0,0047
policy rule coefficient	0,2642	0,2365	0,2029	0,1913
output variability	0,0084	0,0087	0,0076	0,0077
inflation variability	0,0022	0,0025	0,0059	0,0056
convergence speed	a_P: δ = 0.4193		a_P: δ = 0.1616	
	α: δ = 0.4670		α: δ = 0.4525	
	ψ: δ = 0.4765		ψ: δ = 0.4918	
β =5.667				
mean a_P	0,0000	-0,0003	0,0009	0,0014
policy rule coefficient	0,2642	0,2365	0,0975	0,0805
output variability	0,0084	0,0087	0,0086	0,0105
inflation variability	0,0022	0,0025	0,0017	0,0026
convergence speed	a_P: δ = 0.4193		a_P: δ = 0.3331	
	α: δ = 0.4670		α: δ = 0.4544	
	ψ: δ = 0.4765		ψ: δ = 0.4870	

welfare loss
modest (<10%)

speed high but
< 1/2

Some simulation results (8)

- An additional experiment is conducted assuming the economy is hit by a sequence of 12 (gradually declining) shocks. With RE, the impact is short-lived, while under imperfect knowledge the adjustment of the PLM to the ALM acts as a propagation/amplification mechanism.
- As shown in Table 5a, activist policies do not seem to pay off: the lower β, the higher the variability and the wider the fluctuations of both inflation and output.
- Given the structure of the model, the similarities with the OW paper (2003) are striking:
  - 1. the model has no intrinsic dynamics; accordingly, the uncoupling between actual and perceived inflation is much harder to achieve;
  - 2. though output expectations do not play any role, output gap uncertainty affects the accuracy of the central bank's estimates of the moments of α.
- The justification for the existence of a bias in favour of inflation-averse policies is to be found in the role of central bank learning: too activist policies reduce the information content of the output gap and make the estimates of the policy rule too unreliable.

Some simulation results (9)

Table 5a - Dynamic response to contractionary shocks
(decreasing gain sequence)

	Lexicographic preferences		Quadratic preferences	
	UE	*CE*	*UE*	*CE*
β = 0.176				
$\Sigma(y-y^{RE})^2$	0,0437	0,0401	0,0447	0,0484
min y	-0,0177	-0,0149	-0,0132	-0,0196
max y	0,0025	0,0025	0,0067	0,0067
$\Sigma(\pi-\pi^{RE})^2$	0,0115	0,0031	0,0101	0,0090
min π	-0,0028	-0,0012	0,0286	0,0261
max π	0,0013	0,0013	0,0330	0,0330
β = 1.0				
$\Sigma(y-y^{RE})^2$	0,0437	0,0401	0,0449	0,0453
min y	-0,0177	-0,0149	-0,0138	-0,0144
max y	0,0025	0,0025	0,0054	0,0054
$\Sigma(\pi-\pi^{RE})^2$	0,0115	0,0031	0,0049	0,0049
min π	-0,0028	-0,0012	0,0048	0,0048
max π	0,0013	0,0013	0,0084	0,0084
β =5.667				
$\Sigma(y-y^{RE})^2$	0,0437	0,0401	0,0411	0,0398
min y	-0,0177	-0,0149	-0,0162	-0,0149
max y	0,0025	0,0025	0,0028	0,0028
$\Sigma(\pi-\pi^{RE})^2$	0,0115	0,0031	0,0047	0,0017
min π	-0,0028	-0,0012	-0,0003	0,0008
max π	0,0013	0,0013	0,0025	0,0025

lexicographic preferences worse when β is high

the lower β the wider the fluctuations in π e y

too activist policies educe the information ontent of y and biases the estimate of ϕ

Conclusions

- The paper focuses on two claims, namely that
  - 1. policies that are efficient under RE can perform very poorly when knowledge is incomplete and agents learn adaptively;
  - 2. the optimal degree of monetary policy discretion is obtained with policies that put a cap on inflation.
- The main findings are the following:
  - 1. when agents learn adaptively, the incentives and constraints facing the monetary authority change substantially and a bias against activist policies arises;
  - 2. a policy that involves a cap on inflation is helpful in reducing output and inflation variability but it is not uniformly superior to policies aimed at minimising a quadratic loss function;
  - 3. the existence of a bias against activism and the ranking of monetary policies do not depend on whether agents have finite or infinite memory.
- These findings confirm the OW claim, which is remarkable, since the paper relies on a different model and a more complex learning process.