Comments on Michael Woodford's

"Robustly Optimal Monetary Policy with Near-Rational Expectations"

George W. Evans
University of Oregon

December 2, 2006

Summary

- This paper looks at robustly optimal monetary policy in a purely forward-looking NK model

- Robustness is with respect to deviations of expectations from RE.

- The main findings include:

  Commitment remains important – even more so than under RE.

  The long-run inflation target is unchanged

  History dependence remains important

  The response of optimal policy to cost-push shocks is concave in the amplitude of the shocks.

Overall Comments

- This is an important extension of robustness tools to private-sector expectations.

- The paper may not go far enough in its concerns for robustness against private-sector expectations. Are too many deviations ruled out?

- The paper goes too far in its concerns for robustness against private-sector expectations.
  Put differently: Even in this post-9/11 world, I am concerned about evil-agent paranoia. Is this the way to run (economic) policy?

- Robustness vs. Learning

Two-minute Summary

We consider the (linearized) model

$$\pi_t = \kappa x_t + \beta \hat{E}_t \pi_{t+1} + u_t, \tag{2.1}$$

where $\hat{E}_t \pi_{t+1}$ are distorted PS (private sector) beliefs, and $u_t = \vartheta' s_t$, where s_t is an exogenous process

$$s_{t+1} = A s_t + B w_{t+1} \text{ where } w_t \sim N(0, I).$$

The CB (central bank) controls π_t subject to the constraint (2.1). It knows the s_t process and chooses and commits to a policy $\pi(h_t)$, where $h_t = (w_t, w_{t-1}, \ldots)$. Their objective is to minimize

$$E_0 \sum_{t=0}^{\infty} \beta^t \frac{1}{2} \left[\pi_t^2 + \lambda \left(x_t - x^* \right)^2 \right].$$

PS beliefs about s_{t+1} are absolutely continuous with respect to (the correct) CB beliefs. Hence

$$\hat{E}_t \pi_{t+1} = E_t m_{t+1} \pi_{t+1} \text{ for some } m_{t+1} \geq 0 \text{ a.s. with } E_t m_{t+1} = 1.$$

The set-up has a Stackelberg form. The CB picks the policy rule $\pi(h_t)$. Given $\pi(h_t)$ the evil PS agent chooses $m_{t+1}(w_{t+1})$ subject to $E_t m_{t+1} = 1$ to

$$\max_{m_{t+1}(w_{t+1})} E_0 \left\{ \sum_{t=0}^{\infty} \beta^t \left[\frac{1}{2} \left(\pi_t^2 + \lambda \left(x_t - x^* \right)^2 \right) - \theta E_t m_{t+1} \log(m_{t+1}) \right] \right\}. \tag{2.4}$$

$\theta > 0$ measures the penalty to the evil PS agent of a distortion m_{t+1}.

The CB knows it is faced with an evil agent and chooses a robust policy $\pi(h_t)$ to minimize the maximized value of (2.4).

MW solves the problem for linear policies

$$\pi_{t+1}(w_{t+1}) = p_t^0 + p_t^{1\prime} w_{t+1},$$

assuming also the timeless perspective. The coefficients take the form

$$\begin{aligned} p_t^1 &= \bar{p}^1 \text{ and} \\ p_t^0 &= \mu p_{t-1}^0 + a' s_t + \mu \bar{p}^{1\prime} w_t \end{aligned}$$

MW computes details in the univariate white noise case $u_t = \sigma_u w_t$ and compares the results with RE, obtaining the qualitative results mentioned earlier.

The technical apparatus is impressive and the approach will likely generalize.

One question before my main comments: usually a fully robust policy is infeasible. Here MW illustrates results using quite low values of θ. It would be of interest to know if there is a well-defined policy for the limiting "fully robust" case $\theta \longrightarrow 0$.

Comments

1. The paper does not go far enough in examining deviations from RE.

It is appealing to have robustness against unspecified deviations from RE. But is the set of deviations too restrictive? These take a very specific form. PS knows the optimal policy rule

$$\pi_{t+1}(w_{t+1}) = p_t^0 + \bar{p}^{1\prime} w_{t+1},$$

and knows the coefficients. The deviations m_{t+1} concern the law of motion for the exogenous innovation w_{t+1} since

$$\hat{E}_t \pi_{t+1} = p_t^0 + \bar{p}^{1\prime} E_t m_{t+1} w_{t+1}.$$

This seems very restrictive.

In particular:

(a) The absolute continuity assumption means that the PS cannot believe that something might happen if in fact it cannot.

(b) Many small distortions are ruled out, e.g. a small shift right of the perceived density of w_{t+1} or small misperceptions by PS in the values of p_t^0 and $\bar{p}^1$.

(c) The formulation seems to rule out natural LS learning rules. Suppose PS posits a law of motion

$$\pi_{t+1} = \alpha + \sum_{i=0}^{K} \delta_i s_{t-i} + \eta_{t+1},$$

estimates the coefficients by LS (least squares), and uses the estimated equation to forecast inflation, updating the coefficients each period. I suspect this is not an allowed distortion, but surely we want to worry about deviations like these.

After all:

(i) if LS learning converges to RE then it is near-rational and we would surely want to include it as an allowed deviation.

(ii) if LS learning does not converge, or is sometimes far from RE, then we may want to worry about it even more.

(d) More generally, suppose agents include in their regressions extraneous variable (sunspots) or omit key variables. Wouldn't we want policy-makers to be concerned with these?

2. Does the paper goes too far in its concerns for robustness against private-sector expectations?

Within the allowed class of deviations from RE the CB insists on always guarding against the worst outcomes in terms of PS expectations.

(a) It seems to make more sense for the CB to make some assessment of the most plausible deviations of expectations from RE and use a subjective probability distribution over these to find the optimal policy.

(b) The CB may be able to learn over time how PS forms expectations. For example, if PS uses a LS learning rule to form expectations, then the CB might be able to determine this.

(c) Even restricting attention to the limited class of deviations from RE considered, PS would be able over time to estimate the distribution of w_{t+1}. CB policy should reflect this (and the CB could facilitate it by telling PS the true distribution).

(d) More generally the point is that the objective of agents is to maximize utility and they therefore have an incentive to forecast accurately. Agents will not, period after period, choose forecasts simply to embarrass policymakers. Do we want to base policy on the assumption that the PS is malevolent. (And if they really are evil, why would they constrain themselves by $E_t m_{t+1} = 1$?)

I understand that the evil agent assumption is a fiction, used to guard against worst-outcomes, but it remains an extreme assumption.

3. Robustness vs. learning.

- The LS learning literature posits particular forms of deviation from RE by the PS (and/or the CB) in which coefficients of a forecast function are updated over time using LS. As MW stresses, it would be unwise for the CB in practice to devise policies that depend heavily on the assumption that PS follows a particular learning rule. His paper shows how you can design policy when the expectation mechanism may be any member of a specified class.

- A major strand of the papers on monetary policy and learning, however, is that one should design policy so that it guides the PS under learning towards RE (and even towards the RE optimal policy). That link is dropped in the current paper.

- My view is that robustness and learning should be combined. For example, I have some work in which the CB is uncertain about key structural parameters, and optimizes policy subject to this uncertainty and also to ensuring that the resulting REE is stable under LS learning by the PS.

- MW's paper argues that policy should be robust against a general class of deviations from RE. This is clearly desirable, but the current paper drops the feedback from policy to expectations. An alternative formulation worth pursuing would be to look for policy that is (in some sense) robust against a large class of learning rules by the PS, e.g. against adaptive expectations, constant gain LS, LS with inertia and random gains, underparameterized LS, overparameterized LS, SG learning, etc.

Conclusions

- The contribution of the paper is that it considers in a general way the policy problem when PS expectations can deviate from RE within a specified class. This is an important problem.

- MW shows how to solve for robust policy in the minimax sense. This is a technical achievement.

- I remain unconvinced that guarding against evil PS agents is a sound basis for policy.

- Future work in this area should extend the class of deviations from RE and integrate PS learning with policy robustness.