INTERNATIONAL RESEARCH FORUM ON MONETARY POLICY 2006

Discussion of 'Bayesian and Adaptive Optimal Policy under Model Uncertainty' by Lars Svensson and Noah Williams

Eric Schaling

Department of Economics, University of Johannesburg and CentER for Economic Research, Tilburg University, The Netherlands

Svensson and Williams (2006)

- Policymaker seeks to set policy optimally. True structure unobserved and policymakers optimally learn from their observations of the economy

- Classic problem of learning and control, now also forward-looking variables present. Bayesian optimal policy (BOP) numerically solved. Curse of dimensionality: only feasible in relatively small models

- Instead, adaptive optimal policy (AOP) – updating beliefs, but no experimentation. Advantage: easier to compute

- BOP typically includes an experimentation component. In line with e.g. Beck and Wieland (2002), Wieland (2000, 2006), Tesfaselassie, Schaling and Eijffinger (2006).

- AOP provides good approximation to the BOP

- The benefits of learning (moving from NL to AOP) may be substantial, whereas the benefits from experimentation (moving from AOP to BOP) are modest or even insignificant $\Rightarrow$ in *practical monetary policy settings* the experimentation motive may not be a concern

ARE THOSE FINDINGS PLAUSIBLE?

Learning and Control in a Svensson-Type Model

- Tesfaselassie, Schaling and Eijffinger (2006)[1] incorporate forward-looking interest rate expectations in a standard inflation forecast targeting framework, where the true structure is unobserved

- We find that the BOP is in general *more activist*, in the sense of responding aggressively to the state of the economy, than the AOP[2] for small to moderate deviations of the state variable from its target.[3]

- Figure 1 from TSE (2006) shows the response of the control (the interest rate r_t) to deviations of the state w_t from its target level of zero, for a specific belief characterized by the mean $\tilde{c}_0 = 1$ and variance $\tilde{p}_0 = 0.5$.[4]

- Policy reaction functions different, therefore welfare (loss) likely to be different as well

$\Rightarrow$ **BOP $\neg\approx$ AOP (NON-TRIVIAL LOSS DIFFERENCE)!**

AOP provides good approximation to the BOP?

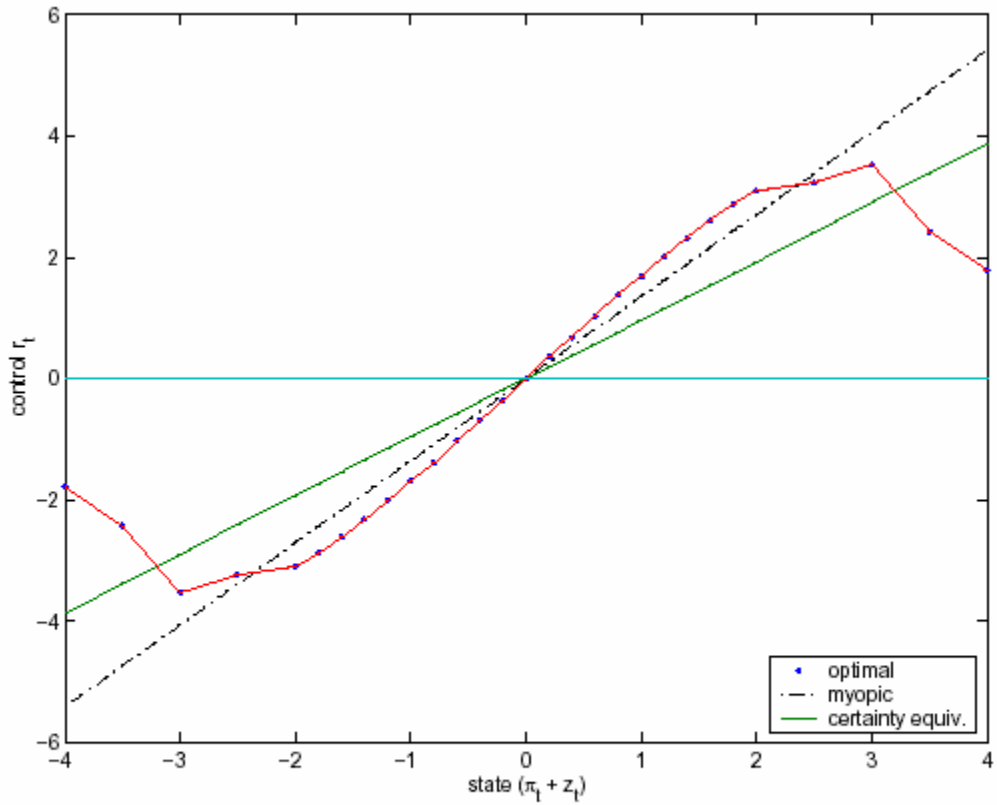

Figure 1: The three decision rules for initial beliefs $\tilde{c}_0 = 1$ and $\tilde{p}_0 = 0.5$

- SW (2006): the benefits of learning (moving from NL to AOP) may be substantial, whereas the benefits from experimentation (moving from AOP to BOP) are modest or even insignificant ⇒ in *practical monetary policy settings* the experimentation motive may not be a concern

- SW (2006): if this preliminary finding stands up to scrutiny, experimentation in economic policy in general and monetary policy in particular may not be very beneficial

DOES THIS FINDING STAND UP TO SCRUTINY?

Table 1 Model Aspects that Positively Affect the Experimentation Motive ($L(AOP) - L(BOP)$)

	Svensson-Williams (2006)	**Tesfaselassie, Schaling and Eijffinger (2006)**
Large uncertainty about true state	? for SW	Yes
Small deviation of state from target	? for SW	Yes
Model becomes more forward-looking	Ambiguous (6.3)[5]	Yes[6]
Shocks driving the data process become less volatile.	? for SW	Yes
Larger weight on stabilizing the control	No (6.1)[7],? for SW (6.3)[8]	Yes[9]
Policy cares more about the future (larger discount factor)	? for SW	Yes[10]

Table 2 Model Aspects Used to Scrutinize Svensson and Williams

	Svensson-Williams (2006)	**Tesfaselassie, Schaling and Eijffinger (2006)**
Large uncertainty about true state	? for SW	Yes
Small deviation of state from target		
Model becomes more forward-looking		
Shocks driving the data process become less volatile.	? for SW	Yes
Larger weight on stabilizing the control		
Policy cares more about the future (larger discount factor)		

The Practical Importance of Experimentation in Monetary Policy

- As the DGP *approaches* a non-stochastic process in TSE (2006) - very low σ_v^2 - and large uncertainty about true state - large $\tilde{p}_0$ (combination of rows 1 and 3) $\Rightarrow$ very difficult to learn true state without experimentation[11] $\Rightarrow$ marginal benefit of experimentation huge $\Rightarrow$ large welfare difference between BOP and AOL policies

- If DGP exhibits high variance then ceteris paribus the opposite is true

Conclusions

- SW (2006): the benefits of learning (moving from NL to AOP) may be substantial, whereas the benefits from experimentation (moving from AOP to BOP) are modest or even insignificant ⇒ in *practical monetary policy settings* experimentation in monetary policy may not be very beneficial

- ES: this statement is CORRECT, but I have shown that this does appear to be a result that heavily depends on (i) the state (what do we know, and where are we), and (ii) on specific parameter configurations

[1] See Tesfaselassie, M., Schaling, E. and Eijffinger, S, 'Learning About the Term Structure and Optimal Rules for Inflation Targeting', CEPR Discussion Paper, No 5896, October 2006.

[2] In TSE (2006) the AOP is essentially the case of passive learning without updating; i.e. certainty equivalence.

[3] For *large* deviations, the BOP is *less activist* than the AOP. The intuition for this is the following. From the updating equations, the larger the deviation of w_t from zero (due to say an exogenous shock), the smaller $\tilde{p}_{t+1}$ (implying that $\tilde{c}_{t+1}$ is a more precise estimate) leading to a smaller control error when setting r_{t+1}. Thus, in contrast to the myopic policy, the actively learning central bank anticipates future improvements in policy performance as $|w_t|$ increases. This shows that when realized exogenous shocks, v_t, that ultimately drive data generation for w_t are large, there is less of a role for deliberate probing by the central bank, more so the larger the deviation of w from the target. This shows that the role of optimal policy in generating variability (increasing the coefficient of variation of the target variable) depends on the current state of the economy. When next period's state of the economy is expected to deviate a lot from the target due to an unpredictable shock, and thus generates data by itself, optimal policy does not need to increase the coefficient of variation of the next period's state w_{t+1}, while it does so when the economy is hit by a very small shock and the state is close to the target.

[4] SW (2006) use X_t and i_t for the control. Same variables on axes (state on horizontal, control on vertical axis, cfr. Their Figure 6.2.c).

[5] The effects of the forward-looking constraint are summarized by varying Ξ_t. When $\Xi_t < 0$, the experimentation component is greater for positive values of X_t and smaller for negative values. The converse happens when $\Xi_t > 0$, as then the experimentation component is smaller for positive values.

[6] As the model becomes more forward-looking, the long term interest rate will be determined to a larger extent by movements in private sector expectations. Due to this fact, the importance of learning increases as well.

[7] Here the period loss function satisfies $L_t = \frac{1}{2}X_t^2.$, which – unlike the standard dual control specification – does not contain the control variable (vector). Put differently, there is no control cost term. Thus, a comparative analysis with respect to that weight cannot be performed.

[8] Here the loss function *does* contain the control, as it obeys $L_t = \frac{1}{2}x_t^2 + 0.1 i_t^2, \quad \delta = 0.95.$.

[9] The incentive to deviate from the AOP diminishes for a central bank that gives more attention to inflation (state) stabilization. In the limit of strict inflation targeting optimal policy is not affected by uncertainty about the true state.

[10] The benefits from experimenting with the policy rate are in terms of reduced variability of the economy in the future from more precise estimates and improved control associated with less uncertainty in the unknown parameter. Under flexible inflation targeting, optimal policy exploits the tradeoff between these *future benefits* and *current costs* from large movements in r_t. In this regard, one would expect that the incentive to experiment increases as the central bank gives more weight to expected future losses from variability in w_{t+1} and r_{t+1} , for $w_{t+\tau}$ and $r_{t+\tau}$ for $\tau = 1,2,3...$ Thus there is more tendency to probe for a relatively large discount factor.

[11] The intuition is that, with low variability in the shocks, optimal policy needs to actively manage data generation by increasing the (conditional) coefficient of variation in w_{t+1}, defined here as $CV = \sqrt{F_{t+1}} / E_t w_{t+1}$. Since F_{t+1} is predetermined, the CV can be increased only by a lower value of $E_t w_{t+1}$, which in turn requires more aggressive response of r_t to w_t .