Chapter 9

Monetary policy rules

9.1 Motivation

So far in the course the focus has been on constructing stochastic dynamic general equilibrium models which match various characteristics observed in the data. To some extent this has been successfully achieved so we now proceed to consider how monetary policy itself should be designed. We are able to do this because the models that have been developed are derived from fairly solid microfoundations, in other words we do not fall foul of the Lucas critique. In this lecture we assess different rules for setting monetary policy.

9.2 Key readings

Rotemberg and Woodford (1998) "Interest-Rate Rules in an Estimated Sticky Price Model" *NBER Working Paper*, No. 6618

Taylor (1993) "Discretion Versus Policy Rules in Practice", *Carnegie-Rochester Conference Series on Public Policy*, 39, 195-214

9.3 Related reading

Chow (1980) *Analysis and Control of Dynamic Economic Systems*, Wiley

Hansen and Sargent (2001) "Robust Control and Model Uncertainty", *mimeo*

Levin, Wieland and Williams (1998) "Robustness of Simple Monetary Policy Rules under Model Uncertainty", *NBER Working Paper*, No. 6570

Rotemberg and Woodford (1998) "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy: Expanded Version", *NBER Technical Working Paper*, No. 233

Tetlow and von zur Muehlen (1999) "Simplicity Versus Optimality: the choice of monetary policy rules when agents must learn", *Board of Governors of the Federal Reserve System Finance and Economics Discussion Series*, No. 1999-10

9.4 An optimising framework

The paper by Rotemberg and Woodford (1998) pulls together many of the theoretical issues we have studied so far. At the heart of their framework is a dynamic IS curve based on intertemporal maximisation and an aggregate supply curve based on the sticky prices in the New Keynesian Phillips curve. Rather than assuming a cash-in-advance constraint and facing the problems of generating substantial liquidity effects, they jump directly to a formulation in which the instrument of monetary policy is the interest rate itself. The central bank is assumed to set interest rates as a function of past rates, current and past inflation and current and past values of the output gap. In other words, interest rates are set according to a rule of the form in equation (9.1).

$$r_t = r^* + \sum_{k=1}^{n_r} \mu_k \left(r_{t-k} - r^*\right) + \sum_{k=0}^{n_\pi} \phi_k \left(\pi_{t-k} - \pi^*\right) + \sum_{k=0}^{n_y} \theta_k \left(y_{t-k} - y^*\right) \tag{9.1}$$

In this equation the central bank is constrained to react linearly to current and past economic variables. However, as we will see later this is not a binding constraint because rules such as (9.1) (subject to including a sufficient number of lags) encompass the optimal monetary policy.

Rotemberg and Woodford (1998) have a neat way of calibrating their model. They wish to calibrate/estimate their theoretical model to match the variances of the three endogenous variables: interest rates, inflation and output. The trick is to recognise that these variances can be completely described by (i) the variances of the fundamental disturbances and (ii) the impulse responses of the three variables to the fundamental disturbances. Hence, instead of matching variances directly, they match the impulse response functions of the structural model to those estimated using a simple structural vector autoregression. This approach has the considerable advantage that shocks in the model have effects similar to those observed in the data, for example Figure 9.1 shows the estimated and theoretical responses of output to a monetary policy shock.

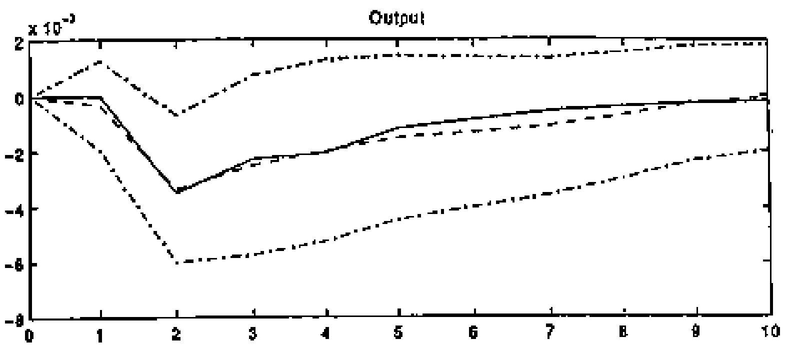

Figure 9.1: Responses of output to a monetary policy shock in the model and data.

The dynamic IS curve, the aggregate supply curve and the interest rate rule (9.1) completely describe the structure of the economy. However, in order to decide on what is the best monetary policy, i.e. what the μ, ϕ and θ coefficients in the monetary policy rule should be, we need to define some target for what monetary policy is trying to achieve. Traditionally, it has been assumed that monetary policy should minimise some quadratic loss function (9.2), where losses are caused by inflation or output being away from their respective targets. χ is the relative weight placed on inflation variability relative to output variability.

$$\mathcal{L} = \sum_{t=0}^{\infty} \beta^t \left[(y_t - y^*)^2 + \chi (\pi_t - \pi^*)^2 \right] \tag{9.2}$$

This is not a completely satisfactory objective for a central bank. If the model is micro-founded on the basis of a representative optimising agent then maximising the welfare of that agent is the appropriate target for central bank policy. In an important innovation, Rotemberg and Woodford (1998) do exactly this and derive the central bank's objective function directly from the welfare of the representative agent. They begin with the utility function (9.3), in which welfare depends on the expected utility of consumption c_t and the expected disutility of the work needed to produce output y_t in each state of the world z.

$$\mathcal{W} = E \left\{ u(c_t) - \int_0^1 v(y_t(z))dz \right\} \tag{9.3}$$

Rotemberg and Woodford (1998) take a second order Taylor series approximation of this to give an objective function of the form (9.4).

$$\mathcal{W} = -\Omega \left(L + {\pi^*}^2 \right)$$

$$L = var(\pi_t) + \psi^{-1} var(\pi_t - E_{t-2}\pi_t) + \Lambda var \left\{ E_{t-1}(\hat{y}_t - y_t^s \right\} \tag{9.4}$$

This is similar in structure to the *ad hoc* quadratic loss function (9.2), although the match is not perfect. If the central bank did minimise (9.2) then there is no guarantee that the welfare of the representative agent (9.4) would be maximised. In this framework the problem of the central bank is to choose the monetary policy which maximises the welfare of the representative agent (9.4) subject to the dynamic IS curve and the New Keynesian aggregate supply curve. The objective (9.4) is quadratic and the constraints can be log-linearised so the problem is in the standard form of linear-quadratic control. For this type of problem, it is possible to show that the optimal monetary policy is indeed of the form shown in Equation (9.1), with interest rates reacting linearly to interest rates, inflation and the output gap. In other words, Equation (9.1) does describe the optimal policy and the problem of the central bank reduces to that of finding the correct coefficients μ, ϕ and θ and the correct lags n_r, n_π and n_y for the monetary policy rule.[1]

9.5 Optimal monetary policy

The full optimal rule as calculated by Rotemberg and Woodford (1998) is shown in Equation (9.5).

[1] For more details of how this is done using the Matrix Riccati equation see Chow (1980), pp. 156-160.

$$
\begin{aligned}
\hat{r}_t \;=\;& 0.29\hat{r}_{t-1} + 0.42\hat{r}_{t-2} + 1.28\hat{r}_{t-3} + 0.22\hat{y}_{t-1} - 0.25\hat{y}_{t-2} + \cdots \\
& +0.16\hat{\pi}_t + 1.00\hat{\pi}_{t-1} + 2.45\hat{\pi}_{t-2} - 1.45\hat{\pi}_{t-3} + 0.74\hat{\pi}_{t-4} \\
& -0.08\hat{\pi}_{t-5} + 0.25\hat{\pi}_{t-6} + 0.33\hat{\pi}_{t-7} + 0.23\hat{\pi}_{t-8} + 0.25\hat{\pi}_{t-9} \\
& +0.19\hat{\pi}_{t-10} + 0.17\hat{\pi}_{t-11} + 0.13\hat{\pi}_{t-12} + 0.19\hat{\pi}_{t-13} + 0.06\hat{\pi}_{t-14} + \cdots
\end{aligned}
\tag{9.5}
$$

The omitted terms in $\hat{y}_{t-j}$ are all of size smaller than 0.01 and the omitted terms in $\hat{\pi}_{t-j}$ are all of size 0.03 or smaller. If coefficients of 0.0001 were to be included then the optimal policy would need 18 lags of output and 49 lags of inflation. Such complexity raises obvious questions about the operational feasibility of fully optimal monetary policy rules such as (9.5).

9.6 Simple monetary policy rules

In contrast to the complex monetary policy rules outlined above, Taylor (1993) suggested a very simple rule for setting interest rates (9.6).

$$i_t = 0.5y_t + 1.5\pi_t \tag{9.6}$$

This paper has been enormously influential, generating a whole industry of research. As Taylor himself comments, it is simple enough to put on the back of a business card. Interest rates are set according to the current output gap and inflation. Clearly such a rule is not optimal so why would a central bank ever want to follow this rule? One answer is that it may be a good approximation to the optimal rule. It is certainly a fair approximation to the behaviour of the Federal Reserve Board over the past forty years. Figure 9.2 shows the path of the actual Federal funds rate (the key instrument for monetary policy in the US) and the rate that would have prevailed if monetary policy had followed the Taylor rule. The two series match each other quite well, especially bearing in mind the considerable difficulties involved in obtaining a reasonable estimate of the current output gap. This evidence suggests that the Taylor rule is a good approximation of the behaviour of the Fed. Whether that implies optimality depends on whether the Fed is believed to have acted optimally over the period.

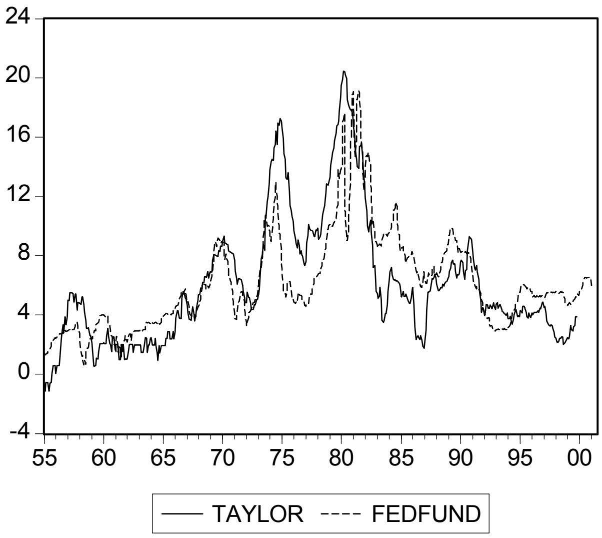

Figure 9.2: The Federal funds rate and Taylor rule based interest rate.

In addition to matching the behaviour of the Federal Reserve Board, simple rules have several other desirable features. We will consider two possibilities: robustness and transparency.

9.7 Robustness

In optimal control theory, policies are distinguished between the good, the bad and the optimal. The complex rule calculated for the Rotemberg-Woodford economy is an example of an optimal rule, the question remains whether it is a good or bad rule. In the context of the Rotemberg-Woodford framework it is a good rule - in fact the best by definition. However, if the economy is not well characterised by the Rotemberg-Woodford model (maybe the economy is not New Keynesian after all) then the optimal policy could turn out to be a very bad policy. What the policy maker would really like is a policy rule which performs well across a wide range of models and in a variety of situations. Such a rule would be robust to errors made in the specification of the underlying model. A recent paper by Levin, Wieland and Williams (1998) assesses whether the Taylor rule may be a good rule in this sense. They analyse the performance of a modified Taylor rule against more complex rules across three different underlying models (the Taylor Multi-Country model (TMCM), the MSR model of Orphanides and Wieland, and the FRB staff model). Their first step is to calculate the optimal weights for a modified Taylor rule in the FRB model, resulting in the rule (9.7).

$$r_t = r_{t-1} + 0.8(\pi_t - \pi^*) + 1.0y_t \tag{9.7}$$

This rule is then applied to each model in turn and compared to what could have been achieved if the Taylor rule had been fine-tuned to each model. The results are shown in Table 9.1.

Model	Loss with Taylor rule fine-tuned to model	Loss with Taylor rule fine-tuned to FRB model	Difference
TMCM	3.61	3.64	0.03
MSR	0.33	0.33	0.00
FRB	2.02	2.02	0.00

Table 9.1: Robustness of Taylor rules

The results indicate that the Taylor rule gives a robust performance in the face of specification error. Even the worst-case scenario, a Taylor rule tuned to the FRB model when the correct model is the Taylor Multi-Country model (TMCM), does not lead to large welfare losses. We now turn our attention to the extent to which more complex rules are similarly robust to model misspecification. The computational burden of calculating such rules is considerable so Levin, Wieland and Williams (1998) restrict themselves to 12-parameter or 8-parameter rules. To give an example, they calculate a 12-parameter rule optimised for the MSR model and apply it to both the FRB and TMCM models. The results are shown in Table 9.2.

Model	Loss with optimal rule fine-tuned to MSR model
TMCM	3.78
MSR	0.16
FRB	1.92

Table 9.2: Robustness of optimal rules

Fine-tuning the monetary policy rule to the MSR model leads to a significant reduction in loss if the MSR model is the correct specification. The loss falls from 0.33 in Table 9.1 to 0.16 in Table 9.2. However, if there is misspecification and the TMCM is really the correct model then this optimal rule performs worse than the Taylor rule. The loss in Table 9.1 under the FRB Taylor rule was only 3.64 but the loss with the optimal MSR rule rises to 3.78 in Table 9.2. In this sense the optimal policy is not robust to model misspecification. It is bad rather than good.

The analysis of Levin, Wieland and Williams (1998) has been criticised because the models it considers are all rather similar. In such an analysis it is not surprising that generally *all* monetary policy rules are quite robust. A more satisfactory approach is to formally search for robust policies across well-defined misspecifications. According to robust control theory, the best policy is a min-max strategy, which effectively entails choosing the policy action which gives the least costly (i.e. leads to the smallest deviation of inflation from target) worst outcome, given an assumption about the range of misspecifications that may occur. Recent work by Hansen and Sargent (2001) and others has developed these ideas more fully.

9.8 Transparency

One further advantage of a simple monetary policy rule is that it is easily understood. Private agents are able to quickly learn the policy of the central bank, an issue examined by Tetlow and von zur Muehlen (1999). They consider rules of the type (9.8) in a Fuhrer-Moore type model.

$$r_t = rr^* + E_{t-1}\pi_t + \beta_R(R_{t-1} - \pi_{t-1}) + \beta_\pi(\pi_{t-1} - \pi^*) + \beta_y y_{t-1} + u_t \tag{9.8}$$

A simple Taylor rule is equal to equation (9.8) with the added restrictions $\beta_R = 0$ and $\beta_\pi = \beta_y = 0.5$. This rule is then compared with optimal 2-parameter rules (with $\beta_R = 0$) and optimal 3-parameter rules. In steady-state there is full information and rational expectations so private agents know the parameters of the rule in force. For this case the expected losses are given in Table 9.3, where losses have been normalised by the loss under the optimal 3-parameter rule.

Rules	Loss
Taylor	1.58
Optimal 2-parameter	1.02
Optimal 3-parameter	1.00

Table 9.3: Steady-state losses under different policy rules

In steady-state it is obvious which rule is the best: the optimal rule is optimal. Now, though, consider a central banker who inherits one of the policy rules. Is it profitable to switch to a new policy rule which gives better steady-state performance? The answer depends on how quickly private agents are able to learn the new coefficients in the policy rule and the costs incurred during the learning process. Learning is not instantaneous since there is always some random noise u_t in the policy rule which disguises the policy maker's true intended actions. Tetlow and von zur Muehlen (1999) report simulation evidence for the welfare cost of switching between different monetary policy rules. Table 9.4 highlights some of their results.

Switch	Percentage change in loss
Taylor ⇒ Optimal 2-parameter	- 18 %
Taylor ⇒ Optimal 3-parameter	- 13 %
Optimal 2-parameter⇒ Optimal 3-parameter	+ 4 %

Table 9.4: Welfare consequences of changing the monetary policy rule

The simulations show that it is always beneficial to switch away from the Taylor rule. However, there are larger gains in terms of reduced welfare loss in switching to an optimal 2-parameter rule than an optimal 3-parameter rule. This is because the 2-parameter rule is easier for private agents to learn. Even though in steady-state the 3-parameter rule does better, the costs of introducing it (due to private agents having to go through a protracted learning process) are higher than for introducing the 2-parameter rule. Similarly, if the central bank inherits a 2-parameter rule then there are no incentives to switch to the 3-parameter rule, even though it is better in the long run. Such results may explain why simple rules have gained so much favour.