Special Feature A

Optimal Control In The Monetary Model Of Singapore[1]

Introduction

In the last three decades, central banks have focused primarily on achieving price stability, while balancing other policy objectives such as output and employment stabilisation. Optimal monetary policy was seen as a choice of how best to manage the short-run trade-off between these goals while ensuring that the long-run objective of price stability was met. This endeavour has been reflected in the way monetary policy is incorporated into macroeconomic models, either by specifying a target path for the policy instrument or through simple feedback rules, such as the well-known Taylor rule.

Another method that has stood the test of time, and gained renewed prominence recently, is optimal control policy.[2] Optimal control techniques are regularly deployed in both the physical and social sciences to solve for the trajectories of control variables in dynamic systems, in order to achieve pre-specified goals. In the monetary policy arena, the optimal control approach entails solving a large-scale macroeconomic model to find the path of the policy instrument that would achieve macroeconomic stability. In this case, the goals are typically the minimisation of deviations of inflation and unemployment from their respective targets.

The aim of this Special Feature is to describe and illustrate the use of an optimal control facility that was recently added to the Monetary Model of Singapore (MMS). Launched in 2000, the MMS is the flagship model used by the Economic Policy Group (EPG) at MAS for the purpose of monetary policy analysis.[3] The model is routinely used to generate economic forecasts, conduct scenario analysis, and perform policy simulations. Hence, the incorporation of an optimal control feature into the MMS is part of EPG's ongoing efforts to enhance its quantitative economic toolkit.

The feature begins with a succinct exposition of the optimal control methodology introduced into the MMS and relates it to the discussion of loss functions in the central banking literature. The optimal control policy is then applied in a retrospective historical setting, and the implied macroeconomic consequences are compared with actual outcomes. Finally, the sensitivity of the results to alternative prioritisations of policy objectives is examined.

---

1 This feature was done in collaboration with Christopher Murphy, Director of Independent Economics and a Visiting Fellow at the Australian National University. Mr Murphy is a consultant to EPG, MAS.

2 Janet Yellen gave prominence to the Federal Reserve's use of optimal control methods in her 2012 speech, when she compared an optimal control policy path with the Taylor rule, and showed that the optimal path would only raise the Federal Funds Rate around early 2016 in order to lower unemployment more quickly and allow inflation to overshoot its target for some time. In a subsequent speech in October 2016, Yellen suggested that hysteresis—the adverse impact on the supply-side of the economy due to persistent shortfalls in aggregate demand—could potentially be reversed by temporarily running a "high-pressure economy" with robust aggregate demand and a tight labour market. See Yellen (2012), Yellen (2016) and Brayton, Laubach and Reifschneider (2014).

3 MAS (2014a) provides a description of EPG's suite of models.

The Optimal Control Methodology

The MMS is a dynamic Computable General Equilibrium (CGE) model which explicitly accounts for the interrelationships between the supply and demand sides of the economy. It recognises that the demand side is important in influencing economic activity in the short run, and is therefore grounded in the New Keynesian tradition. At the same time, the model converges to a neoclassical steady-state growth path dictated by supply-side constraints in the long run. In terms of model structure, the MMS is split into separate equation blocks for domestic demand, trade, the labour market, and sector-specific production functions. In addition, the model encapsulates the impact of fiscal policy, which is assumed to be exogenously given, while the monetary policy instrument in Singapore—the S$NEER—serves to anchor the paths of prices and other nominal variables in the model.[4] Although monetary policy has effects on real economic activity in the short to medium run, it is neutral in the long run.

The optimal control solution embedded in the MMS seeks to attain given macroeconomic goals, while minimising short-term changes in the monetary policy instrument. Specifically, the algorithm solves for the paths of the S$NEER and the fiscal policy instrument that minimise the costs of current and future deviations of the inflation and unemployment rates from their target values, as represented by the following loss function:

$$L=\sum_{t=1}^{T}\left(\frac{1}{1+\delta}\right)^{t}\left\{\begin{array}{l}\alpha_1\left(\pi_t-\overline{\pi}\right)^2+\alpha_2\left(u_t-\overline{u}\right)^2\\+\alpha_3\left[\Delta\left(r_t-r_t^*\right)\right]^2+\alpha_4\left(s_t-\overline{s}\right)^2\\+\alpha_5\left(\tau_t-\tau_{t-4}\right)^2\end{array}\right\}$$

where π denotes CPI inflation, u the seasonally adjusted resident unemployment rate, r and r^* represent the domestic and foreign interest rates, respectively, s denotes the public sector surplus as a ratio of GDP, τ denotes the effective income tax rate and Δ is the first difference. Long-run target values are denoted with a bar above.

The first two components of the loss function impose a penalty on the squared deviations of inflation and unemployment rates from their target values. For illustrative purposes, the application presented in the next section assumes the target for the resident unemployment rate to be around 3.5%—the long-run average since the early 2000s. Likewise, the target inflation rate is taken to be the average rate of CPI inflation, of 1.8%, in the last three decades.

The third argument of the loss function imposes a cost on instrument instability, i.e., abrupt changes in the monetary policy instrument. This cost is captured by the squared difference of the interest rate differential, and aims to minimise the spillovers from exchange rate movements on domestic monetary conditions through the uncovered interest parity condition. However, the penalty on the change in the exchange rate applies only after the first forecast period. This allows for a discrete change in the exchange rate in the first period, if needed, thus conferring considerable policy flexibility to the optimal control method.

The last two components in the loss function pertain to fiscal policy. The fourth term ensures that the government's intertemporal budget constraint is adhered to in the long run while the last term penalises large fluctuations in the fiscal instrument, which is assumed to be the income tax rate in the MMS. Since the loss function is specified over the entire forecast horizon from time t to T, a discount factor δ that places a larger weight on nearer periods has been added.[5]

The α parameters are the relative weights on each argument of the loss function. Terms with higher weights will be more strictly binding in the optimal control exercise. The weight on the inflation target, α_1, is normalised to unity and all other weights are calibrated relative to α_1. The weight on the resident unemployment target is set at 2 because the unemployment rate is

4 See MAS (2014b) for detailed information on the MMS.

5 A typical real social discount rate of approximately 5% per annum is used.

historically less volatile than inflation, and therefore deviations from the target should be penalised more heavily. The weights on the instrument stability terms are relatively small, but sufficient to avoid erratic swings in the policy instruments. The sensitivity of the optimal policy path to different weights is discussed below.

The solution to the minimisation problem described above involves an iterative procedure implemented within the MMS that solves for the present and future values of the control variables. Let x be defined as the vector of the time paths of the control variables (i.e., policy instruments) and $y = h(x)$ the vector of the time paths of the target variables $\{\pi, u, \tau, \Delta(r - r^*), s\}$, subject to the dynamic interdependencies between the exogenous and endogenous variables in the MMS.

If y^* is taken to denote the vector of the desired paths of y, the optimal control solution is the value of x that minimises the quadratic loss function given by:

$$f = \left(y - y^*\right)' W \left(y - y^*\right)$$

where W is a diagonal matrix of the weights α_i, $i = 1, \ldots, 5$ for every time period.

In principle, the solution to the optimal control problem makes use of the first derivatives (gradient) of the loss function as well as the matrix of second derivatives (Hessian). Operationally, the optimal control solution is arrived at through an iterative procedure typically involving several rounds of MMS simulations.

Application Of Optimal Control In The MMS

In this section, the optimal control feature in the MMS is used to derive a hypothetical path for the S$NEER over the time period 2007–13.[6] This period is selected as it covers the key global events that had a strong bearing on inflation outcomes in Singapore—the food and energy price shocks in 2007–08, as well as the Global Financial Crisis (GFC) and its aftermath. During this period, inflation rose from 2.1% to a peak of 6.6% in 2008, then fell and troughed at 0.6% in 2009. It increased again over 2010–12, but eased back to 2.4% in 2013.

Some caveats should be borne in mind in interpreting the optimal control policy paths derived from the MMS. First, the optimal control solution will always result in greater macro stability by design, relative to actual policy. Second, the optimal solution for the policy instrument is typically made conditional on the baseline set of economic forecasts and on the target values chosen. In the historical illustration that follows, however, the optimal path of the S$NEER is derived on an *ex-post* basis i.e., it is conditional on the actual outcomes. In other words, perfect foresight of all shocks that had occurred is assumed and the optimal policy path represents the ideal trajectory if policymakers adhered strictly to their loss function and had been able to fully anticipate economic outcomes. Third, monetary policy under optimal control is more activist than actual policy responses, as it is allowed to change every quarter, given the frequency of the model data set-up. In practice, monetary policy is announced biannually and off-cycle moves are undertaken only under rare circumstances, when the policy planning parameters shift abruptly and significantly.

Fourth, the inflation and unemployment rate outcomes under the optimal control approach do not immediately achieve their targets because of inherent lags in policy transmission, as well as the penalties imposed on instrument instability.

Historical Simulation Results

Charts 1(a)–(d) present the actual path of the S$NEER and realised outcomes for three macroeconomic variables, namely, headline CPI inflation, the seasonally adjusted resident

[6] Given the set-up of the optimal control problem, a solution for the fiscal instrument is also obtained in this exercise. However, the discussion in this feature will be confined to monetary policy.

unemployment rate and the real GDP growth rate. In the same chart, the model-generated outcomes from the optimal control policy are also shown.

The actual S$NEER values, in both direction and magnitude, are relatively close to the theoretically derived values from the model simulation results. Indeed, the S$NEER path obtained from the optimal control approach generally lies within ±1% of the actual S$NEER. In this regard, it is important to note that actual movements in the S$NEER would also reflect fluctuations within the prescribed policy band, whereas the optimal control solution confines the exchange rate to a specific path.

Nonetheless, it would be instructive to examine more closely the episodes during which the actual S$NEER path deviated relatively more significantly from the optimal control solution. These periods were in 2007, H1 2008 and Q4 2009 – Q4 2010. The differences between the actual outcomes and the optimal control solutions are noteworthy as they reflect the other important considerations that MAS took into account in policymaking.

In the first episode, while MAS tightened policy at the end of 2007, this came later and by less than that suggested by the optimal control solution. This was due in part to the price effects associated with the impending GST hike in July 2007. In particular, due consideration was given to the one-off impact on consumer prices of the tax change, as well as the buffer provided by other offsetting fiscal measures, which would have tempered the price impact on the real disposable incomes of households. In contrast, the optimal policy solution prescribes an immediate tightening as it targets the overall CPI inflation rate, which would have captured the full impact of the GST hike. The initial upward jump in the exchange rate along the optimal policy path reflects the fact that the design of the loss function does not penalise movements in the first forecast period. Moreover, the dampened policy response by MAS reflected the conscious decision to accommodate the uncertainty arising from the US subprime crisis at the time.

Optimal policy also indicated a more accommodative policy path than what transpired in April 2008. In this case, MAS did not ease policy by as much in the face of rising global commodity price pressures, compounded by tight labour market conditions and escalating property prices domestically. The pertinent consideration here was the possible interaction of several sources of shocks to inflation—both supply and demand factors, in addition to asset market dynamics—which was assessed to warrant a more pre-emptive approach, given the possible upside risks to short-term inflation expectations.

In the third episode following the GFC, the smaller-than-prescribed tightening of the monetary policy stance was a measured move, given the still tentative recovery from the crisis. In this instance, the optimal control path over the next few quarters had not fully accommodated the significant downside risks to the baseline growth and inflation outcomes prevailing at the time, and hence the need to adopt a more cautious approach towards tightening the policy stance.

The differences between the actual and optimal policy paths illustrate the additional considerations that impinge on the monetary policy formulation process, which cannot be made fully endogenous in a model simulation. While the optimal control results are informative, they tend to overlook some factors affecting the growth-inflation trade-off as well as the flexibility of Singapore’s exchange rate-based monetary policy framework. These include the nature and source of shocks, which could cause greater variability in inflation. A supply-side shock for example, could be short-lived and also induce an optimising adjustment to spending behaviour in response to (relative) price movements. Further, the presence of uncertainty over the near-term baseline path for the economy adds another important dimension to policy formulation. Under uncertainty, policy would need to take into account the asymmetric risks and costs that could arise from growth and inflation deviations.

Chart 1
Comparison of Actual Vs Optimal Control

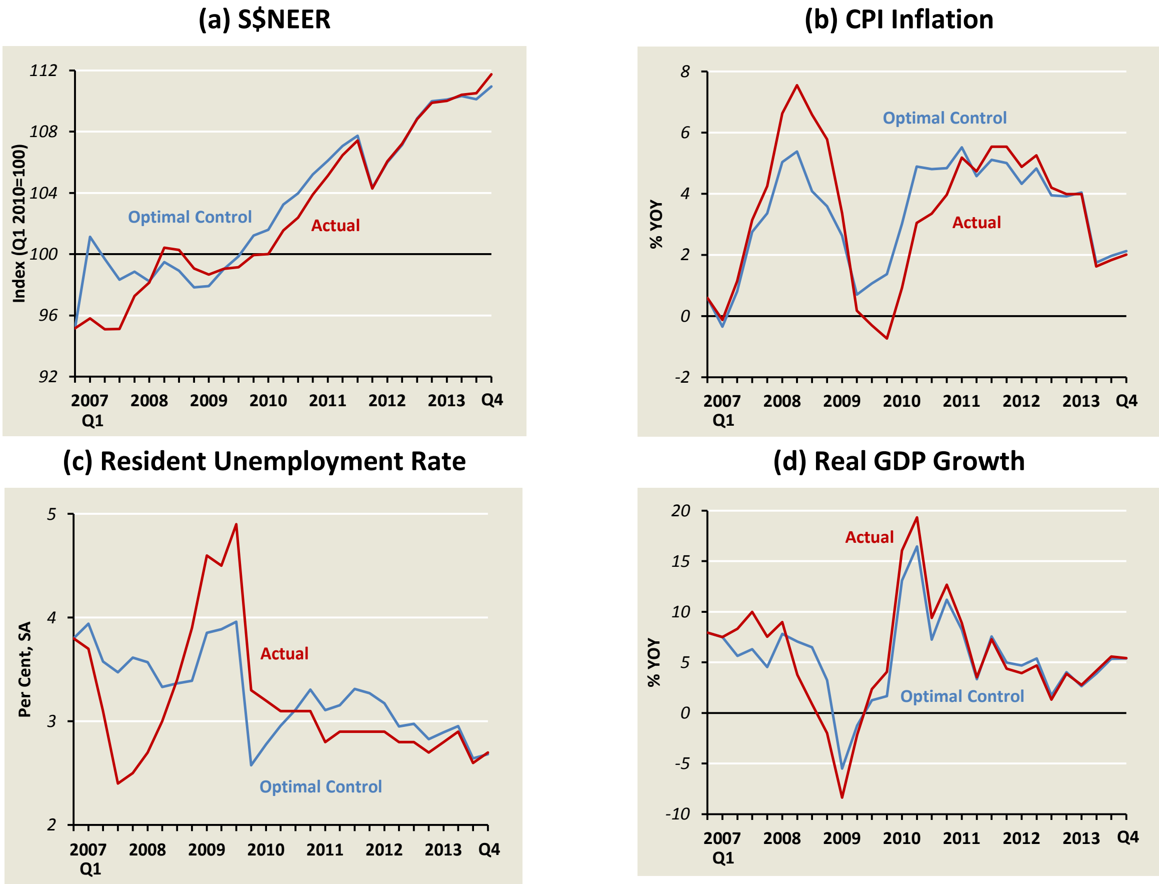

Alternative Loss Function Specifications

This section examines how the optimal control policy and macroeconomic outcomes change when the relative weights on the policymaker's loss function are altered. To establish upper and lower bounds around the optimal path shown earlier, two markedly different pairs of relative weights are assigned to the inflation and unemployment rate arguments in the function. In the first experiment, the weights assigned to inflation and unemployment deviations are raised from $(\alpha_1 = 1,\ \alpha_2 = 2)$ to $(\alpha_1 = 1,\ \alpha_2 = 10)$. Accordingly, this places a much larger weight on the costs of unemployment. In the second simulation, the weights are changed to $(\alpha_1 = 20,\ \alpha_2 = 2)$, thus switching the policy emphasis to keeping inflation closer to its long-run norm.

Charts 2(a)–(d) plot the optimal S$NEER path and associated macroeconomic outcomes from these two pairs of relative weights. Owing to the short-run Phillips curve trade-off between unemployment and inflation, the inflation rate converges to its long-run target at a faster pace when a larger weight is placed on the inflation deviation term in the loss function. However, this comes at the expense of higher unemployment and output volatility. The converse is true in the case of a larger weight on unemployment, although GDP growth is not much affected.

Across the different specifications, inflation variability tends to be higher than output and unemployment variability. This finding stems from the greater sensitivity of inflation outcomes to changes in the S$NEER, as compared to the other two variables. In addition, there is greater variance in the S$NEER optimal path associated with a larger inflation weight, as the exchange rate would have to be adjusted by more to dampen deviations in the inflation rate. Nonetheless, except for the period of the GFC, the optimal S$NEER paths for the different loss functions generally move in tandem.

Chart 2
Comparison of Optimal Control under Different Weight Specifications

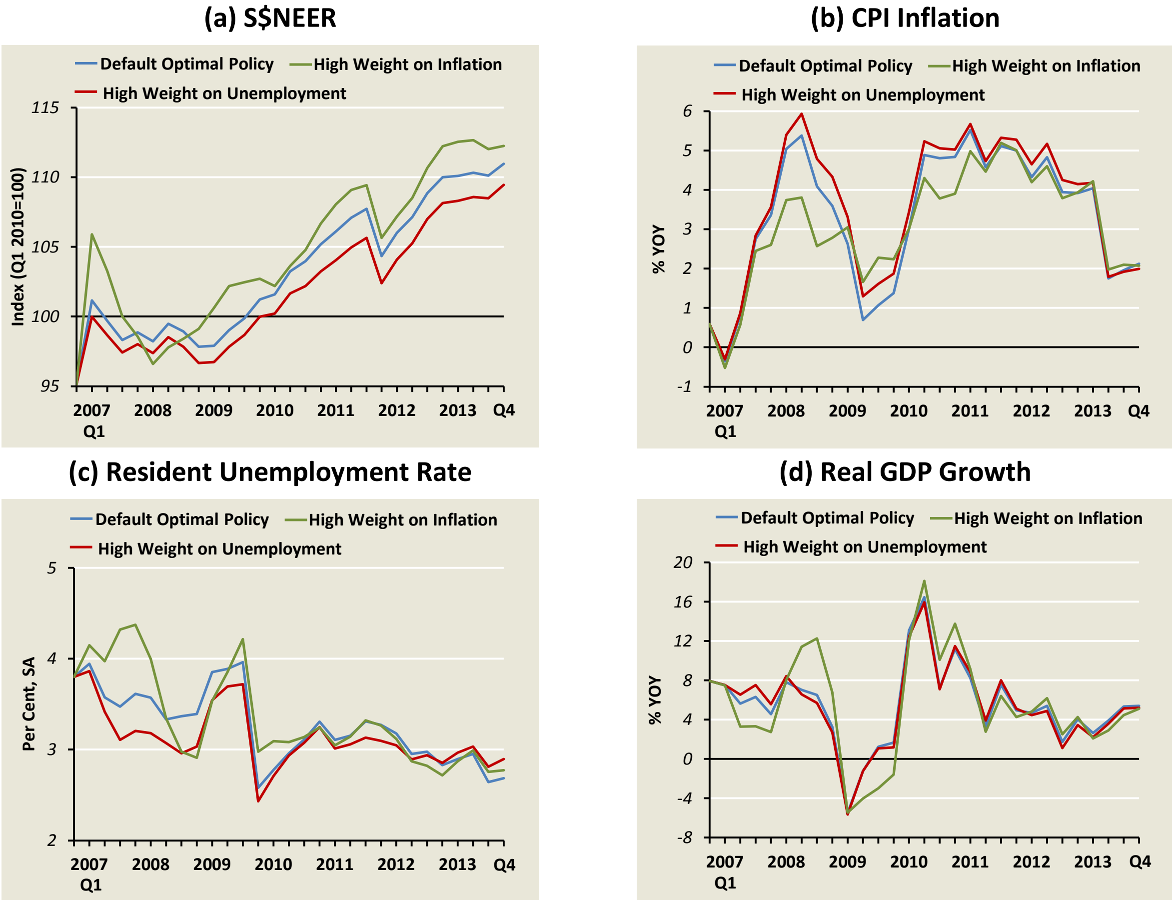

Conclusion

The introduction of an optimal control monetary policy within the MMS provides in principle a systematic and disciplined approach towards arriving at a benchmark against which policy options can be evaluated. Nonetheless, MAS' approach to policy formulation already indirectly incorporates elements of the optimal policy approach through the simulation of alternative policy paths and evaluation of the resultant macroeconomic outcomes.

As described in this feature, the application of optimal control on a historical baseline demonstrates both the usefulness and limitations of the method. Specifically, the optimal control path is always contingent on the assumed set of macroeconomic forecasts. In this regard, deviations between the optimal and actual paths of the S$NEER can be partly attributed to the assumption of perfect foresight built into the exercise. Moreover, the optimal policy trajectory depends on the policymaker's assumed loss function. Therefore, the optimal control solution should be interpreted with caution and does not represent MAS' *de facto* policy.

As no model can fully capture the workings of the economy, MAS is not wedded to any single model or method to inform monetary policy. In particular, the presence of uncertainty and less-than-perfect knowledge of the economy calls for a considerable degree of judgement in the conduct of monetary policy. Still, the optimal control methodology serves as a useful reference point for practical policy formulation.

References

Brayton, F, Laubach, T, and Reifschneider, D (2014), "Optimal-Control Monetary Policy in the FRB/US Model", *FEDS Notes*.

Monetary Authority of Singapore (2014a), "MAS Macroeconomic Modelling Workshop 2014", *Macroeconomic Review*, Vol. XIII(2), pp. 76–85.

Monetary Authority of Singapore (2014b), "The Monetary Model of Singapore (MMS): A Technical Overview" available at http://www.mas.gov.sg/~/media/MAS/Monetary%20Policy%20and%20Economics/Education%20and%20Research/Education/Macroeconometric%20Models/The%20Monetary%20Model%20of%20Singapore%20MMS%20A%20Technical%20Overview.pdf

Yellen, J L (2012), "Perspectives on Monetary Policy", speech at the Boston Economic Club Dinner available at https://www.federalreserve.gov/newsevents/speech/yellen20120606a.htm

Yellen, J L (2016), "Macroeconomic Research after the Crisis", speech at the 60th annual economic conference sponsored by the Federal Reserve Bank of Boston available at https://www.federalreserve.gov/newsevents/speech/yellen20161014a.htm