

Stochastic Control of Econometric Models for Slovenia

Dimitri Blueschke, Viktoria Blueschke-Nikolaeva, and Reinhard Neck^(✉)

Department of Economics, Alpen-Adria-Universität Klagenfurt,
Universitätsstraße 65-67, 9020 Klagenfurt, Austria
`reinhard.neck@uni-klu.ac.at`

Abstract. This paper considers the optimal control of two small stochastic models of the Slovenian economy applying the OPTCON algorithm. OPTCON determines approximate numerical solutions to optimum control problems for nonlinear stochastic systems and is particularly applicable to econometric models. We compare the results of applying the OPTCON2 version of the algorithm to the nonlinear model SLOVNL and the linear model SLOVL. The results for both models are similar, with open-loop feedback controls giving better results on average but with more ‘outliers’ than open-loop controls.

Keywords: Optimal control · Stochastic control · Algorithms · Policy applications · Nonlinear models

1 Introduction

Optimum control theory has been applied in many areas of science, from engineering to economics. An algorithm that provides (approximate) solutions to optimum control problems for nonlinear dynamic systems with different kinds of stochastics is OPTCON, which was first introduced in [3]. An extension has been developed in [2], which includes passive learning or open-loop feedback control policies.

OPTCON was implemented in MATLAB and can deliver numerical solutions to problems with real economic data. Two such applications are described and analyzed in this paper. We develop two macroeconomic models of the Slovenian economy, a nonlinear model called SLOVNL and a (comparable) linear model called SLOVL. The algorithm with both open-loop and open-loop feedback strategies is applied to these models and the influence of the control scheme and the nonlinearity of the model on the optimum solution is investigated in some optimization experiments.

2 The Problem

The OPTCON algorithm is designed to provide approximate solutions to optimum control problems with a quadratic objective function (a loss function to

be minimized) and a nonlinear multivariate discrete-time dynamic system under additive and parameter uncertainties. The intertemporal objective function is formulated in quadratic tracking form, which is quite often used in applications of optimum control theory to econometric models. It can be written as

$$J = E \left[\sum_{t=1}^T L_t(x_t, u_t) \right], \quad (1)$$

with

$$L_t(x_t, u_t) = \frac{1}{2} \begin{pmatrix} x_t - \tilde{x}_t \\ u_t - \tilde{u}_t \end{pmatrix}' W_t \begin{pmatrix} x_t - \tilde{x}_t \\ u_t - \tilde{u}_t \end{pmatrix}. \quad (2)$$

x_t is an n -dimensional vector of state variables that describes the state of the economic system at any point in time t . u_t is an m -dimensional vector of control variables, $\tilde{x}_t \in R^n$ and $\tilde{u}_t \in R^m$ are given ‘ideal’ (desired, target) levels of the state and control variables respectively. T denotes the terminal time period of the finite planning horizon. W_t is an $((n+m) \times (n+m))$ matrix, specifying the relative weights of the state and control variables in the objective function. W_t (or W) is symmetric.

The dynamic system of nonlinear difference equations has the form

$$x_t = f(x_{t-1}, x_t, u_t, \theta, z_t) + \varepsilon_t, \quad t = 1, \dots, T, \quad (3)$$

where θ is a p -dimensional vector of parameters the values of which are assumed to be constant but unknown to the decision maker (parameter uncertainty), z_t denotes an l -dimensional vector of non-controlled exogenous variables, and ε_t is an n -dimensional vector of additive disturbances (system error). θ and ε_t are assumed to be independent random vectors with expectations $\hat{\theta}$ and O_n respectively and covariance matrices $\Sigma^{\theta\theta}$ and $\Sigma^{\varepsilon\varepsilon}$ respectively. f is a vector-valued function fulfilling some differentiability assumptions, $f^i(\dots)$, is the i -th component of $f(\dots)$, $i = 1, \dots, n$.

3 The Optimum Control Algorithm

The OPTCON1 algorithm [3] determines policies belonging to the class of open-loop controls. It either ignores the stochastics of the system altogether or assumes the stochastics to be given once and for all at the beginning of the planning horizon. The nonlinearity problem is tackled iteratively, starting with a tentative path of state and control variables. The tentative path of the control variables is given for the first iteration. In order to find the corresponding tentative path for the state variables, the nonlinear system is solved numerically. After the tentative path is found, the iterative approximation of the optimal solution starts. The solution is sought from one time path to another until the algorithm converges or the maximal number of iterations is reached. During this search the system

is linearized around the previous iteration's result as a tentative path and the problem is solved for the resulting time-varying linearized system. The criterion for convergence demands that the difference between the values of current and previous iterations be smaller than a pre-specified number. The approximately optimal solution of the problem for the linearized system is then used as the tentative path for the next iteration, starting off the procedure all over again.

The more recent version OPTCON2 [2] incorporates both open-loop and open-loop feedback (passive-learning) controls. The idea of passive learning corresponds to actual practice in applied econometrics: at the end of each time period, the model builder (the control agent) observes what has happened, that is, the current values of state variables, and uses this information to re-estimate the model and hence improve his/her knowledge of the system.

The passive-learning strategy implies observing current information and using it in order to adjust the optimization procedure. For the purpose of comparing open-loop and open-loop feedback results, it is not possible to observe current and true values, so one has to resort to Monte Carlo simulations. Large numbers of random time paths for the additive and multiplicative errors are generated, representing what new information could look like in reality. In this way 'quasi-real' observations are created and both types of controls, open-loop and passive-learning (open-loop feedback), are compared.

4 The SLOVNL Model

We estimated two simple macroeconomic models for Slovenia, one nonlinear (SLOVNL) and one linear (SLOVL). The SLOVNL model (**SLO**Venian model, **Non-Linear** version) is a small nonlinear econometric model of the Slovenian economy consisting of 8 equations, 4 behavioral equations and 4 identities. SLOVNL includes 8 state variables, 3 control variables, 4 exogenous non-controlled variables and 16 unknown (estimated) parameters. We used quarterly data for the time periods 1995:1 to 2006:4; this data base with 48 observations admits a full-information maximum likelihood (FIML) estimation of the expected values and the covariance matrices for the parameters and the system errors. The starting period for the optimization is 2004:1; the terminal period is 2006:4 (12 periods).

Model variables used in SLOVNL:

Endogenous (state) variables:

$x[1]$	<i>CR</i>	real private consumption
$x[2]$	<i>INVR</i>	real investment
$x[3]$	<i>IMPR</i>	real imports of goods and services
$x[4]$	<i>STIRLN</i>	short term interest rate
$x[5]$	<i>GDP</i>	real gross domestic product
$x[6]$	<i>VR</i>	real total aggregate demand
$x[7]$	<i>PV</i>	general price level
$x[8]$	<i>Pi4</i>	rate of inflation

Control variables:

$u[1]$	<i>TaxRate</i>	net tax rate
$u[2]$	<i>GR</i>	real public consumption
$u[3]$	<i>M3N</i>	money stock, nominal

Exogenous non-controlled variables:

$z[1]$	<i>EXR</i>	real exports of goods and services
$z[2]$	<i>IMPDEF</i>	import price level
$z[3]$	<i>GDPDEF</i>	domestic price level
$z[4]$	<i>SITEUR</i>	nominal exchange rate SIT/EUR

Model equations:

Standard deviations are given in brackets.

$$\begin{aligned}
 CR_t = & 240.9398 + 0.740333 CR_{t-1} + 0.111727 GDPR_t \left(1 - \frac{TaxRate_t}{100}\right) \\
 & (189.7449) \quad (0.1115) \quad (0.0330) \\
 & - 1.007353 (STIRLN_t - Pi4_t) - 4.773533 Pi4_t \\
 & (2.5848) \quad (2.4966)
 \end{aligned}$$

$$\begin{aligned}
 INVR_t = & 75.41731 + 0.932211 INVR_{t-1} + 0.264523 (VR_t - VR_{t-1}) \\
 & (176.8549) \quad (0.1423) \quad (0.0924) \\
 & - 0.455511 (STIRLN_t - Pi4_t) - 2.981241 Pi4_t \\
 & (6.9044) \quad (3.1277)
 \end{aligned}$$

$$\begin{aligned}
 IMPR_t = & IMPR_{t-1} + 0.826449 (VR_t - VR_{t-1}) - 38.14954 SITEUR_t \\
 & (0.0724) \quad (18.9336)
 \end{aligned}$$

$$\begin{aligned}
 STIRLN_t = & 0.811606 STIRLN_{t-1} - 0.000877 \frac{(M3N)_t}{PV_t} \cdot 100 \\
 & (0.1375) \quad (0.0008) \\
 & + 0.002746 GDPR_t \\
 & (0.0026)
 \end{aligned}$$

$$GDPR_t = CR_t + INVR_t + GR_t + EXR_t - IMPR_t$$

$$VR_t = GDPR_t + IMPR_t$$

$$PV_t = \frac{GDPR_t}{VR_t} \cdot GDPDEF_t + \frac{IMPR_t}{VR_t} \cdot IMPDEF_t$$

$$Pi4_t = \frac{PV_t - PV_{t-4}}{PV_{t-4}} \cdot 100$$

The objective function penalizes deviations of objective variables from their ‘ideal’ (desired, target) values. The ‘ideal’ values of the state and control variables (\tilde{x}_t and \tilde{u}_t respectively) are chosen as shown in Table 1. The ‘ideal’ values for most variables are defined in terms of growth rates (denoted by % in Table 1) starting from the last given observation (2003:4). For $Pi4$ and $TaxRate$, constant ‘ideal’ values are used; for $STIRLN$, a linear decrease of 0.25 per quarter is assumed to be the goal.

Table 1. ‘Ideal’ values of objective variables, SLOVNL

<i>CR</i>	<i>INVR</i>	<i>IMPR</i>	<i>STIRLN</i>	<i>GDPR</i>	<i>VR</i>	<i>PV</i>	<i>Pi4</i>	<i>TaxRate</i>	<i>GR</i>	<i>M3N</i>
1%	1%	2%	-0.25	1%	1.5%	0.75%	3	25.2	1%	1.75%

The weights for the variables, i.e. the constant matrix W in the objective function, are first chosen as shown in Table 2a (‘raw’ weights) to reflect the relative importance of the respective variable in the (hypothetical) policy maker’s

Table 2. Weights of objective variables, SLOVNL

2a: 'raw' weights		2b: 'correct' weights	
variable	weight	variable	weight
<i>CR</i>	1	<i>CR</i>	3.457677
<i>INVR</i>	1	<i>INVR</i>	12.16323
<i>IMPR</i>	1	<i>IMPR</i>	1.869532
<i>STIRLN</i>	1	<i>STIRLN</i>	216403.9
<i>GDPR</i>	2	<i>GDPR</i>	2
<i>VR</i>	1	<i>VR</i>	0.333598
<i>PV</i>	1	<i>PV</i>	423.9907
<i>Pi4</i>	0	<i>Pi4</i>	0
<i>TaxRate</i>	2	<i>TaxRate</i>	37770.76
<i>GR</i>	2	<i>GR</i>	63.77052
<i>M3N</i>	2	<i>M3N</i>	0.090549

objective function. These 'raw' weights have to be scaled or normalized according to the levels of the respective variables to make the weights comparable. The normalized ('correct') weights are shown in Table 2b.

5 The SLOVL Model

To analyse the impact of the nonlinearity of the system we developed a linear pendant to the SLOVNL model. This 'sister model' is called SLOVL (**SLO**venian model, **L**inear version) and consists of 6 equations, 4 of them behavioral and 2 identities. The model includes 6 state variables, 3 exogenous non-controlled variables, 3 control variables, and 15 unknown (estimated) parameters. We used the same data base as for SLOVNL and a specification as close as possible to that of SLOVNL in order to make comparisons between the results of the algorithm for a linear and a nonlinear model. Again, we used full-information maximum likelihood (FIML) to estimate the expected values and the covariance matrices for the parameters and the system errors. The starting and the terminal period for the optimization are again 2004:1 and 2006:4.

Model variables used in SLOVL:

Endogenous (state) variables:

$x[1]$	<i>CR</i>	real private consumption
$x[2]$	<i>INVR</i>	real investment
$x[3]$	<i>IMPR</i>	real imports of goods and services
$x[4]$	<i>STIRLN</i>	short term interest rate
$x[5]$	<i>GDPR</i>	real gross domestic product
$x[6]$	<i>VR</i>	real total aggregate demand

Control variables:

$u[1]$	<i>Taxes</i>	tax revenue
$u[2]$	<i>GR</i>	real public consumption
$u[3]$	<i>M3R</i>	money stock, real

Exogenous non-controlled variables:

$z[1]$	<i>EXR</i>	real exports of goods and services
$z[2]$	<i>SITEUR</i>	nominal exchange rate SIT/EUR
$z[3]$	<i>Pi4</i>	rate of inflation

Model equations:

Standard deviations are given in brackets.

$$CR_t = 231.582776 + 0.744522 CR_{t-1} + 0.111736 (GDPR_t - Taxes_t) - 0.855137 (STIRLN_t - Pi4_t) - 4.657411 Pi4_t$$

(191.99)
(0.11)
(0.03)
(2.63)
(2.50)

$$INVR_t = 69.965212 + 0.936305 INVR_{t-1} + 0.265119 (VR_t - VR_{t-1}) - 0.292918 (STIRLN_t - Pi4_t) - 2.869522 Pi4_t$$

(176.51)
(0.14)
(0.09)
(6.90)
(3.11)

$$IMPR_t = IMPR_{t-1} + 0.826648 (VR_t - VR_{t-1}) - 38.158117 SITEUR_t$$

(0.07)
(18.86)

$$STIRLN_t = 0.811458 STIRLN_{t-1} - 0.000877 (M3R)_t + 0.002748 GDPR_t$$

(0.14)
(0.0008)
(0.0026)

$$GDPR_t = CR_t + INVR_t + GR_t + EXR_t - IMPR_t$$

$$VR_t = GDPR_t + IMPR_t$$

The objective function is analogous as for SLOVNL, where the ‘ideal’ values of the state and control variables (\tilde{x}_t and \tilde{u}_t respectively) are chosen as shown in Table 3. For the weights for the variables, Table 4a shows the ‘raw’ weights and Table 4b gives the normalized weights.

Table 3. ‘Ideal’ values of objective variables, SLOVL

<i>CR</i>	<i>INVR</i>	<i>IMPR</i>	<i>STIRLN</i>	<i>GDPR</i>	<i>VR</i>	<i>Taxes</i>	<i>GR</i>	<i>M3R</i>
1%	1%	2%	-0.25	1%	1.5%	25.2	1%	1.75%

Table 4. Weights of objective variables, SLOVL

4a: ‘raw’ weights		4b: ‘correct’ weights	
variable	weight	variable	weight
<i>CR</i>	1	<i>CR</i>	3.457677
<i>INVR</i>	1	<i>INVR</i>	12.16323
<i>IMPR</i>	1	<i>IMPR</i>	1.869532
<i>STIRLN</i>	1	<i>STIRLN</i>	216403.9
<i>GDPR</i>	2	<i>GDPR</i>	2
<i>VR</i>	1	<i>VR</i>	0.333598
<i>Taxes</i>	2	<i>Taxes</i>	27.52906
<i>GR</i>	2	<i>GR</i>	63.77052
<i>M3R</i>	2	<i>M3R</i>	0.292662

6 Optimization Experiments

The OPTCON2 algorithm is applied to the two econometric models SLOVNL and SLOVL. Two different experiments are run for both models: in experiment 1, two open-loop solutions are compared, a deterministic one where the variances and covariances of the parameters are ignored, and a stochastic one where the estimated parameter covariance matrix is taken into account. In experiment 2, the properties of the open-loop and the open-loop feedback solutions are compared. Furthermore, by comparing the results for the SLOVNL and the SLOVL models we want to analyze the impact of nonlinearity on the properties of the optimal solution.

6.1 Experiment 1: Open-Loop Optimal Policies

For experiment 1, two different open-loop solutions are calculated: a deterministic and a stochastic one. The deterministic solution assumes that all parameters of the model are known with certainty and are equal to the estimated values. In the stochastic case, the covariance matrix of the parameters as estimated by FIML is used but no updating of information occurs during the optimization process.

The results (for details, see [1,2]) show that both the deterministic and the stochastic solutions follow the ‘ideal’ values fairly well but fiscal policies are less expansionary and hence real *GDP* is mostly below its ‘ideal’ values. The values of the objective function show a considerable improvement in system performance obtained by optimization and only moderate costs of uncertainty.

An interesting result is that the deterministic and the stochastic open-loop solutions are very similar. Furthermore, one can see that the SLOVL model is a good ‘linear approximation’ of the nonlinear SLOVNL model because the results for both models are nearly identical. This fact can be used for isolating the impact of nonlinearity on finding the optimum control solution, especially for the case of open-loop feedback policies.

6.2 Experiment 2: Open-Loop Feedback Optimal Policies

The aim of experiment 2 consists in comparing open-loop (OL) and open-loop feedback (OLF) optimal stochastic controls. Figures 1 and 2 show the results of a representative Monte Carlo simulation, displaying the value of the objective function arising from applying OPTCON2 to the SLOVNL and the SLOVL models respectively, under 1000 independent random Monte Carlo runs. The graphs plot the values of the objective function for OL policies (x-axis) and OLF policies (y-axis) against each other. In the ‘zoom in’ panels of the figures, we cut the axes so as to show the mass of the points and omitting ‘outliers’, i.e. results where the value of the objective function becomes extremely large.

One can see that in most cases the values of the objective function for the open-loop feedback solution are smaller than the values of the open-loop solution, indicated by a greater mass of dots below the 45 degree line. This means that




Fig. 1. OL and OLF control, value of objective function; SLOVNL; 1000 Monte Carlo runs; left: ‘normal’, right: ‘zoom in’




Fig. 2. OL and OLF control, value of objective function; SLOVL; 1000 Monte Carlo runs; left: ‘normal’, right: ‘zoom in’

open-loop feedback controls give better results (lower values of the cost function) in the majority of the cases investigated. For the SLOVNL model, the OLF policy gives better results than the OL policy in 66.4% of the cases, for the SLOVL model in 65.4% of the cases considered here.

However, one can also see from these figures (especially in the left-hand panels with a ‘normal’ view) that there are many cases where either control scheme results in very high losses, indicated by dots which are significantly distant from the main mass of the dots. These cases are called ‘outliers’ and can be seen even more clearly in Fig. 3. This figure shows the same results of the 1000 independent Monte Carlo runs for each model (SLOVNL and SLOVL) separately, but for each Monte Carlo run. The OLF and OL objective function values are plotted in Fig. 3 together on the y-axis in each Monte Carlo run, the number of which is shown on the x-axis. Diamonds represent open-loop feedback results and squares represent open-loop results.

The results mean that (passive) learning does not necessarily improve the quality of the final results; it may even worsen them. One reason for this is the presence of the two types of stochastic disturbances: additive (random system




Fig. 3. Open-loop vs. open-loop feedback control, value of objective function (1000 Monte Carlo runs) (left: SLOVNL, right: SLOVL)

error) and multiplicative error (‘structural’ error in the parameters). The decision maker cannot distinguish between realizations of errors in the parameters and in the equations as he just observes the resulting state vector. Based on this information, he learns about the values of the parameter vector, but he may be driven away from the ‘true’ parameter vector due to the presence of the random system error.

6.3 On the Impact of Nonlinearity

In the previous section we saw that there is a severe problem of what we called ‘outliers’ - cases with very high losses or values of the objective function. In a similar framework of optimum stochastic control, [4] investigated numerical reasons for outliers. It was not possible to confirm that the sources of the problem found by these authors were decisive for the outlier problem in our framework. We suspect that there are other reasons for the outliers. One possible reason is the stochastics of the dynamic system itself. In our SLOVNL and SLOVL models all the parameters (including all the intercepts) are considered to be stochastic, which may make this reason more likely to work.

The second possible reason is based on the nonlinear nature of the models for which the OPTCON algorithm was created. The SLOVL model was created mainly in order to test this possibility. The graphical results in the previous section show that the outliers occur in the linear as well as in the nonlinear model version. Moreover, in some of the experiments with 1000 Monte Carlo runs for the SLOVNL model, it turned out that the algorithm did not converge in some runs. In these cases, the algorithm starts to diverge and results in some non-reasonable or even complex numbers for some variables. In the 1000 Monte Carlo runs experiment considered above this happened six times. On the contrary, under the SLOVL model, not one single case of non-convergence out of the 1000 Monte Carlo runs occurred. Thus we arrive at the conclusion that nonlinearity is not the reason for the ‘outliers’, but it can worsen the problem.

7 Conclusion

A comparison of open-loop control (without learning) and open-loop feedback control (with passive learning) shows that open-loop feedback control outperforms open-loop control in the majority of the cases investigated for the two small econometric models of Slovenia. But it suffers from a problem of ‘outliers’ which is present for both policy schemes. When comparing the results for the nonlinear SLOVNL model and the linear SLOVL model, we found that the non-linearity of the system is not responsible for the ‘outliers’ but may worsen their influence in some cases.

References

1. Blueschke, D., Blueschke-Nikolaeva, V., Neck, R.: Stochastic control of linear and nonlinear econometric models: some computational aspects. *Comput. Econ.* **42**, 107–118 (2013)
2. Blueschke-Nikolaeva, V., Blueschke, D., Neck, R.: Optimal control of nonlinear dynamic econometric models: an algorithm and an application. *Comput. Stat. Data Anal.* **56**, 3230–3240 (2012)
3. Matulka, J., Neck, R.: OPTCON: an algorithm for the optimal control of nonlinear stochastic models. *Ann. Oper. Res.* **37**, 375–401 (1992)
4. Tucci, M.P., Kendrick, D.A., Amman, H.M.: The parameter set in an adaptive control monte carlo experiment: some considerations. *J. Econ. Dyn. Control* **34**, 1531–1549 (2010)