Optimal Control of Unemployment in a General Equilibrium Job Search Model

R.D. Herbert[a] and G.D. Leeves[b]

[a]Faculty of Science and Information Technology, The University of Newcastle, Newcastle, Australia.
[b]Department of Economics, The University of Queensland, Brisbane, Australia

Abstract: In this paper optimal control techniques are used for targeting active labor market policies at unemployment in a general equilibrium labour market model that is calibrated to data for the Australian economy in 1998. Observer models are used to develop strategies based on predicted states of the full system. Explicit constraints are introduced that represent realistic limitations on the policy variable based on available government data. The results indicate that observer models improve targeting compared to traditional optimal control methodologies and also provide a better long-term result than a countercyclical policy where policy response is triggered by high unemployment rates.

Keywords: *Labour market policies; Equilibrium search model; Optimal Control.*

1. INTRODUCTION

Two clear facts emerge from recent surveys of the Australian unemployment experience (Dawkins, 2000; Chapman and Kenyon, 2002), which are similar to many other countries. Firstly, the causes of unemployment are complex and interrelated and we do not fully understand how cyclical and structural changes interact and impact on key variables of concern such as unemployment and employment. Secondly, the use of policies to reduce unemployment are often subject to constraints. Chapman and Kenyon (2002) note that countercyclical fiscal policy may be constrained politically by the tax funding implications that it creates. Microeconomic reform may improve efficiency and raise employment in the long run but is likely to increase unemployment in the short-term. The unemployment can be concentrated in particular regions and so may limit the pace at which such reforms are implemented. A recent example in Australia is car tariffs, where the government has slowed the pace of tariff reform (Fisher, 1998). The use of active labour market programs (ALMPS) may be constrained by operational considerations and also the scale effects of the programs. There was a large expansion of wage subsidy programs under the Working Nation initiative (Department of Employment, Education, Training and Youth Affairs, 1996). The required number of placements could not be found in the projected time period and the administrative systems could not place all clients in places that were available.

In this paper we focus on ALMPS as a policy tool to reduce unemployment. The labour market programs are characterised as periods of training that impart a treatment effect to disadvantaged groups (the long-term unemployed) by improving their labour market efficiency. We develop a model where policy makers work in a system that is complex and policy makers have limited knowledge of the operating environment. In addition, policy makers work under specific physical or political constraints. Optimal control methods are employed to develop a model that learns from the full system to derive strategies to achieve desired outcomes given constraints that exist. The application of the procedure involves initially specifying the full model and social loss function both of which are highly non-linear. From this a linear model and quadratic loss function are used to approximate the model and policy objectives. This separate model learns from the behaviour of the system to change and update predictions to achieve targets. It assumes only limited knowledge of the full system. A sliding window of time periods is used for the predictive model to set control policies and these are then updated in the next time period based on the observed states of the system (Herbert and Bell, 2003).

The net benefits of ALMPS remains a key topic of concern (Martin and Grubb, 2001). A recent analysis of the aggregate impact of ALMPS in selected European countries was sceptical about ALMP aggregate unemployment effects (de Koning and Mosley, 2001). The studies were based on partial equilibrium methods, where vacancies and wages were included as exogenous controls. As Martin and Grubb (2001) suggest this is an unrealistic assumption in aggregate analysis. They argue that in the longer term the increase in the effective labour supply produced by programs will generate more vacancies and the higher unemployment associated with an expansion in programs will reduce wage pressures. Hence, we construct a general equilibrium model where wages and vacancies are endogenous and calibrate the model to the Australian labour market.

We use the calibrated model to examine optimal strategies for the targeting of ALMP at long-term unemployment. There have been studies of labour market policies using calibrated general equilibrium models with endogenous wages and vacancies (Gautier and den Butter, 1997) including evaluations of ALMP effects (der Linden and Dor, 2001; Herbert and Leeves, 2003). Typically, policy evaluation is conducted through impulse response analysis by assessing the response of targets to policy shocks. We extend this literature by examining policy evaluation through generating optimal control policies to achieve targets subject to the social loss function of policy makers, whilst introducing explicit constraints on the operation of policies. This develops a policy strategy to meet objectives rather than presenting the results of ad hoc policy changes. We compare the optimal control strategy to a countercyclical response strategy driven by high unemployment rates. Whilst we use ALMPS as the policy variable in the current example the procedure has more general applications for constrained policy analysis.

2. THE LABOUR MARKET MODEL

We have presented the labour market model in earlier research (Herbert and Leeves, 2003; Leeves and Herbert, 2001). There are four labour market stocks in the model; employment, unemployment, vacancies and not-in-the-labour-force. Flows of persons between these stocks are modeled. The model is a nonlinear discrete time model. It consists of 10 equations and has 16 endogenous variables and 5 exogenous variables. It includes equations for wages, vacancies and the matching the unemployed to vacancies. The model distinguishes between short term and long term unemployment. The short term unemployed are model as 12 duration classes with one class for each month of unemployment.

The central feature of the model is a dynamic Cobb-Douglas matching process where a distinction is drawn between short-term (less than a year) and long-term unemployment. The short-term unemployed are more likely to be matched to a vacancy.

In this paper we extend the job matching function to include ALMPs. To operationalise ALMP as a policy, we assume that a proportion (λ) of those entering long-term unemployment at a time period participate in the program for six months, and whilst on the program they cannot engage in any job search activity. After completion of the program they re-enter the long-term unemployed pool but they have had their efficiency increased as a result of the program. The model is calibrated to data for the Australian economy for 1998.

3. ALMP POLICY USING CONTROL

Australian governments have long held out 5% as the target unemployment rate (Committee on Full Employment Opportunities, 1993; Piggott and Chapman, 1995). In 1998 the seasonally adjusted unemployment rate hovered around 7.9%. We adopt as our target a 5% unemployment rate. This is unlikely to be achieved solely through increasing ALMPS, without any growth in worker productivity, changes in distribution of match surplus (which are both exogenous variables in the model) or use of other policy initiatives. Nevertheless, it provides a target that makes our control strategy consistent with the long-term aims of government policy. When we implement constraints on the optimal control they consist of two rules. Firstly, a limit of between 0 and 0.5 for the proportion of long-term unemployed on programs. Clearly, negative proportions are not sensible, but unconstrained optimal control could produce negative values. An upper limit of 50% is chosen even though the highest observed proportion commencing job search training between May 1998 and September 2000 was 0.10. This was a period of a Coalition (conservative) government seeking to restrain ALMP expenditure. During the height of the Working Nation initiative of the Labor government (1995-96) when funds were more available job search training places for long-term unemployed rose to over 0.3. Thus, 0.5 would represent an historically high level of ALMPS associated with a major policy initiative. The baseline value is 0.05.

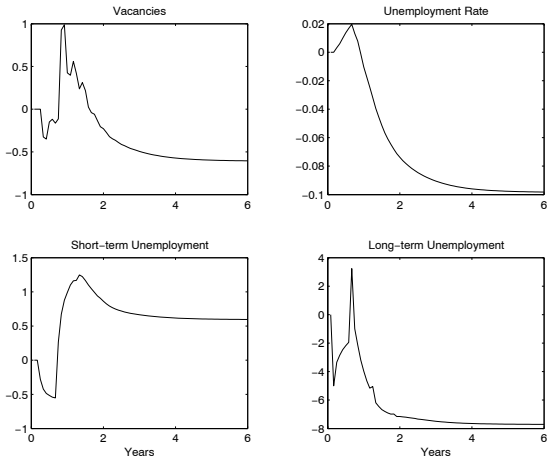

Figure 1: Expansion in Labour Market Programs. (% deviations from baseline.)

3.1. Expansion in ALMPs – No Control

Initially we show the response of the model to a permanent (step) expansion in ALMPs. We raise the value of λ from 0.05 to 0.1. The responses are shown in Figure 1, with the results presented as deviations from the baseline. The unemployment rate falls from its baseline value to 7.78% to 7.68%. It is evident that this reduction comes about entirely from falling numbers of long-term unemployed. Short-term unemployment actually increases, illustrating the spillover effects from the increased numbers of more search efficient long-term unemployed. Vacancies initially rise as firms find it more profitable to post vacancies as they are matched quicker with unemployed searchers. However as employment grows and unemployment falls workers organisations are able to extract a greater proportion of the match surplus reducing the profitability of vacancies. This is indicated by the steady state reduction in vacancies. A permanent increase of ALMPs by 100% helps long-term unemployed and reduces the unemployment rate by 0.1%.

Our next step is to introduce dynamic policy adjustment rules.

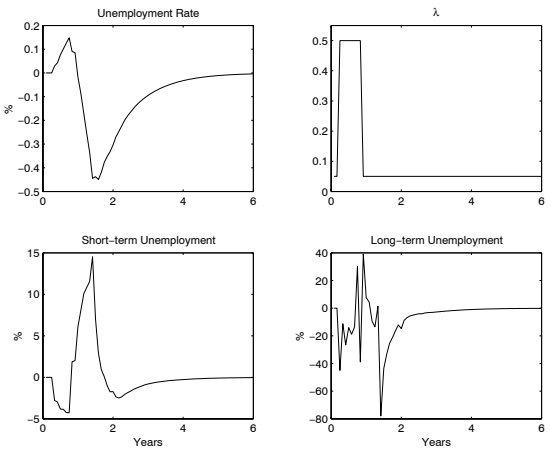

Figure 2: Countercyclical Labour Market Programs. (% deviation from baseline, except λ.)

3.2. Countercyclical ALMPs

As a countercyclical policy we implement a λ of 0.5 whenever the unemployment rate is greater than or equal to its baseline level and a λ of 0.05 (the baseline value) when the unemployment rate is less than the baseline rate. As would be expected from the previous simulation, the simulation of this countercyclical policy will effectively implement an impulse response on the model. The results are shown in Figure 2. The effects of such a policy are short-term, with the longer term effects of the policy being an imperceptible fall in the unemployment rate.

3.3. Optimal Control from Linear Model

As a first optimal control policy, we use the response from a permanent expansion in ALMPs we build a linear model of the form:

$$y_{t+2} = \delta_1 y_{t+1} + \delta_2 y_t + \delta_3 u_{t+1}$$

where y is the output of the linear model, which is the unemployment rate from the nonlinear model; u is the control for the linear model, which is λ in the nonlinear model; and, the δs are the parameters to be estimated. This model was estimated using OLS.

To facilitate the exposition of the development of the various control policy rules, this model is written in state space form as:

$$y_t = Cx_t$$

$$x_{t+1} = Ax_t + Bu_t$$

where

$A=[[0; 1]^T, [\delta_1, \delta_2]^T]^T$, $B=[0; \delta_3]^T$, $C=[0; 1]$ and the states are $x=[y_{t-1}, y_t]^T$.

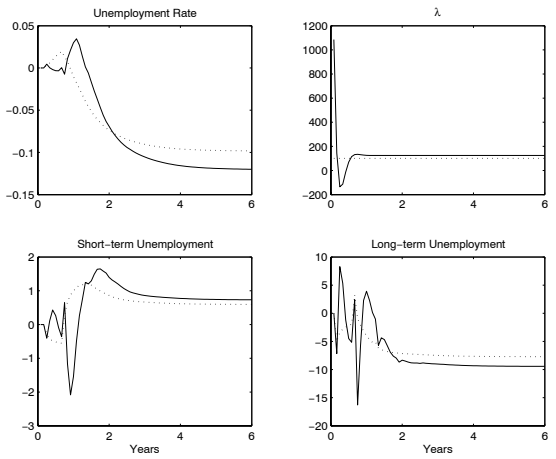

Figure 3: Linear Model Based Control. (% deviation from baseline. Dashed line is the response from Figure 1.)

A policy response (or control) rule is developed on the linear model and applied to the evolution of the variables within the nonlinear model. An optimal control rule was used which minimises the social loss function expressed as deviations from the target output (the unemployment rate). This is a standard linear-quadratic output controller, with a social loss function of:

$$J = \sum_{t=1}^{N} [(y_t - y_{d,t})^T Q(y_t - y_{d,t}) + (u_t - u_{d,t})^T R(u_t - u_{d,t})$$

For the control rule used in the simulations Q is the identity matrix and R = 1. The results are shown in Figure 3. Unemployment is lower in the steady-state than in the previous scenario (as shown by the dashed line in the figure). This control rule introduces substantial variations in λ in the early periods which translates into strong transient fluctuations in short-term and long-term unemployment. One can clearly see that the suggested changes in ALMP commencements fall outside the range of a plausible policy strategy, with the maximum value of λ of over 1000 being nonsensical.

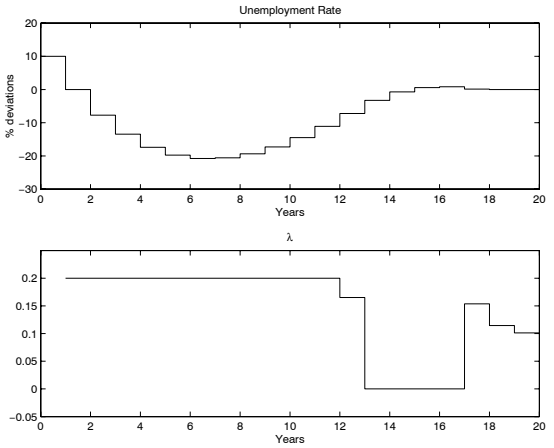

Figure 4: Constrained Linear Model Control.

3.4. Constrained Optimal Control from Linear Model

We continue to develop our policy response from the linear model, but now introduce constrained response. We use the techniques we have presented previously (Herbert and Bell, 2001 and Herbert and Bell, 2003). Basically the technique involves predicting the linear model's output over time and using the observer model as a sliding window on these predictions. The optimal control is implemented through quadratic programming. A subset of the optimal control is then applied to the non-linear model. The whole process is then repeated

As a first step we show the performance of the policy response function on the linear model. In this case we subject the linear model to initial shock (of 10% to the unemployment rate). The aim of the policy is to return the unemployment rate to its pre-shock level. Constraints are placed on the level of λ with $0 \le \lambda \le 0.2$. Neither the shock nor the level of λ are constructed to simulate realistic values but are illustrative. The weighting in the social loss function are set to unity for all time. An eight time-period window is used, and the first control from the window is taken as the policy for that period. The results are shown in Figure 4. The constrained policy response strategy is able to return the model to its baseline. Note that the model is run for a longer time horizon than in previous scenarios. The strategy uses the maximum value of the control (λ) that the rule permits. The transient dynamics show that the strategy creates some overshooting of the target.

When developing the policy response function for the nonlinear model we introduced a learning component to the linear model. The linear model then becomes an observer (Herbert, 1998) for the nonlinear model and adjusts its states to allow for

the difference between the linear and nonlinear models' output. The linear model state equation then becomes:

$$x_{t+1} = Ax_t + Bu_t - L(y_t - Y_t)$$

where Y is the output from the nonlinear model and L is a learning gain. Using the learning model allows the linear model to take into account its prediction error from the nonlinear model through the initial conditions for each sliding time window.

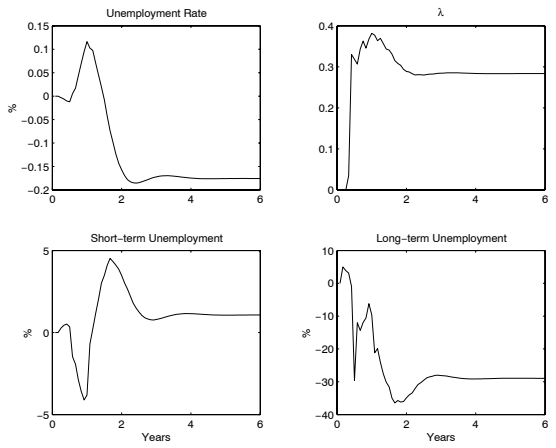

Figure 5: Linear Model Based Predictive Control of Nonlinear Model (A). (% deviation from baseline, except λ.)

The results are shown in Figures 5 and 6. The scenario here is aimed at reducing the unemployment rate. Again the weightings in the social loss function are unity for all time and an eight period sliding window for prediction is used. The learning gain is set to 0.5, so that $L = [0.5,0.5]^T$. Constraints are placed on λ with $0 \leq \lambda \leq 0.5$. The choice of 0.5 is based upon the policy being practically possible but at the upper limit of actual application of ALMPs.

Figure 5 shows that the policy strategy has a value of λ of 0.28 in the steady-state, and results in the unemployment rate falling by 0.2%, long-term unemployment falls substantially and there is a spillover into short-term unemployment. Figure 6 shows that vacancies fall, and wages and employment increase with this policy strategy. The upper constraint on λ, is not reached but the lower constraint is.

An expansion of programs has two offsetting effects. More program participants raise overall efficiency in the matching process thereby increasing outflows from unemployment. However, the resulting increase in employment puts upward pressure on wages, which reduces the outflow of vacancies as their profitability falls. The observer model balances these forces to achieve the best overall outcome for the target unemployment rate. Clearly, the standard optimal control model (Figure3) goes beyond the point where marginal benefits of program exceed the costs of expansion and is wasteful of resources. The observer model has learnt to operate in a more effective manner at a substantially lower cost.

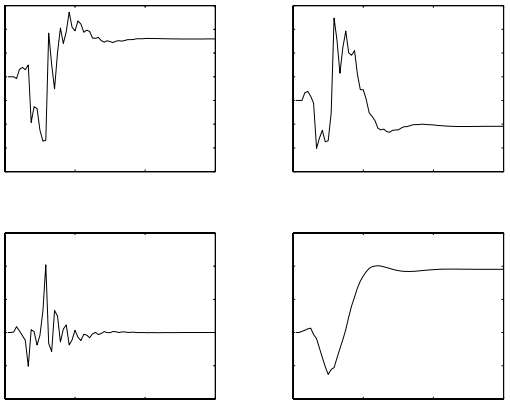

Figure 6: Linear Model Based Predictive Control of Nonlinear Model (B). (% deviation from baseline.)

When compared to Figure 3, where the policy strategy was based on the more usual optimal control approach, we see that the values of λ are realistic, better results are obtained for the unemployment rate and short and long term unemployment; and that there is less oscillatory transient dynamics.

4. CONCLUSION

This paper applied optimal control techniques to the policy variable active labour market programs (ALMPS) targeted at long-term unemployment in the context of a general equilibrium model of the labor market calibrated to reflect the Australian economy in 1998. The target was a reduction in the overall unemployment rate, consistent with Australian government's long-term policy aims. The optimal control was applied using specific policy constraints that reflected limitations in the level and rate at which programs can be expanded or contracted based on available data. Moreover, the optimal control methods used allowed for limited knowledge of the economic system. They were implemented using a separate predictive model that develops a control strategy using a subset of the predicted states, this is referred to as an observer model.

It was found that the observer model was able to achieve better outcomes than standard optimal control techniques in terms of reducing overall

unemployment and the strategy was realistic in terms of operating within the constraints. Moreover, the reduction in long-term unemployment was greater (the policy's immediate target) and spillover effects from increasing the search effectiveness of long-term unemployed on short-term unemployed were slightly reduced. Thus, there was a clear gain from the observer model strategy. The strategy suggested an initial expansion of programs in the first two years near to the constrained limit and then settled to a new higher steady state level of 28% rather than the baseline 5% level. The program effects are understated through no allowance for worker productivity changes and no other complementary policy initiatives have occurred, such as a reduction in replacement ratio. Alternatively, the effects may be overstated if expansion in programs is accompanied by a reduction in average treatment effect due to reduction in training quality or failure in administration support, of which there is some evidence in the Australian case. Nevertheless, the results demonstrate that the optimal strategy consists of a sustained expansion in programs and indicate that a countercyclical strategy for ALMPS is inferior in the long-term. The expansion in employment will raise tax revenue and reduce outlays and this has to be set against the costs of the program to evaluate the budgeting cost.

5. REFERENCES

Chapman, B. and Kenyon, P. (2002), 'Creating jobs-the role of government: Introduction', Australian Journal of Labour Economics 5(2), 101–110.

Committee on Full Employment Opportunities (1993), Restoring Full Employment: A Discussion Paper, Australian Government Publishing Service, Canberra.

Dawkins, P. (2000), The Australian labour market in the 1990s, in D. Gruen and Shrestha, eds, 'The Australian Economy in the 1990s', Reserve Bank of Australia, Sydney.

de Koning, J. and Mosley, H., eds (2001), Labour Market Policy and Unemployment, 1st edn, Edward Elgar.

Department of Employment, Education, Training and Youth Affairs (1996), Working Nation: Evaluation of the Employment, Education and Training Elements, EMB Report 2/96, Australian Government Publishing Service, Canberra. Evaluation and Monitoring Branch.

der Linden, B. V. and Dor, E. (2001), 'Labor market policies and equilibrium employment: Theory and application for Belgium'. Universite Catholique de Louvain, Institut de Recherches Economiques et Sociales, mimeo.

Fisher, T. (1998), Government response to the industry commission report on the automotive industry, Media release, Department of Finance and Trade.

Gautier, P. and den Butter, F. (1997), 'Job creation, wage formation and un-employment dynamics in an empirical flow model for the labour market', pp. 1–23. Free University, Amsterdam, mimeo.

Herbert, R. D. and Bell, R. D. (2003), 'Constrained macroeconomic policy development with a separate predictive model', Mathematics and Computers in Simulation, (forthcoming).

Herbert, R. D. and Leeves, G. D. (2003), 'Labour market policies and long-term unemployment in a flow model of the Australian labour market', Australian Economic Papers. (forthcoming).

Leeves, Gareth D. and Ric D Herbert, (2001), "Unemployment Dynamics in a Flow Model of the Australian Labour Market", *Proceedings of MODSIM2001 – International Congress on Modelling and Simulation,* Canberra, Australia, December 2001, eds. Fereidoun Ghassemi, Michael McAleer, Les Oxley and Michelle Scoccimarro, vol. 3, pp 1415-1420.

Martin, J. P. and Grubb, D. (2001), What works and for whom: a review of OECD countries' experiences with active labour market policies, Working Paper 2001:14, IFAU - Office of Labour Market Policy Evaluation.

Piggott, J. and Chapman, B. (1995), 'Costing the job compact', Economic Record 71(215), 313–328.