DISCUSSION PAPER SERIES

No. 6957

OPTIMAL MONETARY POLICY USING A VAR

Vito Polito and Michael R Wickens

INTERNATIONAL MACROECONOMICS

Centre for Economic Policy Research

www.cepr.org

Available online at: **www.cepr.org/pubs/dps/DP6957.asp**

ISSN 0265-8003

OPTIMAL MONETARY POLICY USING A VAR

Vito Polito, University of York
Michael R Wickens, University of York and CEPR

Discussion Paper No. 6957
September 2008

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 6957

September 2008

ABSTRACT

Optimal Monetary Policy using a VAR*

In this paper we propose a new way to formulate optimal policy based on a quadratic intertemporal welfare function where the dynamic constraint is based on a VAR model of the economy which we call the PVAR method. We argue that the VAR under control should not be derived simply by replacing the VAR equation for the policy instruments by an optimal control rule because this alters the stochastic structure of the VAR. Instead, one should first transform the VAR in order to condition the non-policy variables on the policy instruments, then use the resulting sub-system as the dynamic constraint, and finally construct the VAR under control by combining this sub-system with the resulting optimal policy rule. In this way the original stochastic structure of the VAR is retained. In comparing the two approaches we explain the theoretical advantages of the PVAR over the standard method and we illustrate the methods by examining the formulation of optimal monetary policy for the US. We suggest that since the whole process is easily automated, the PVAR method may provide a useful benchmark for use in real time against which to compare other, probably far more labour intensive, policy choices.

JEL Classification: C2, C6 and E5
Keywords: monetary policy, optimal control and VAR models

Vito Polito
Department of Economics
and Related Studies
University of York
York
YO10 5DD
Email: vp117@york.ac.uk

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=163403

Michael R Wickens
Department of Economics
and Related Studies
University of York
York
YO10 5DD
Email: mrw4@york.ac.uk

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=100329

* The authors would like to thank Richard Baillie, Martin Ellison, Patrick Minford, Adrian Pagan, Hashem Pesaran, Peter Phillips and Mark Salmon for their comments on an earlier draft of the paper.

Submitted 01 September 2008

1 Introduction

Given a welfare function and knowledge of the structure of the economy, monetary policy based on commitment to a rule is usually superior to one based on discretion, yet almost all central banks use discretion. There are several reasons for this. First, most monetary authorities do not have a clearly articulated welfare function. This also holds true for a monetary authority who is formally a strict inflation targeter. Second, even if the monetary authority did have such a welfare function, it is not clear how the monetary policy rule should be specified. A Taylor (1993) rule is commonly invoked, but it is unlikely to be optimal. In order to obtain an optimal rule it is necessary to maximise the welfare function subject to a model of the economy. Third, it is often difficult to get agreement on how to specify a structural model of the economy. In these circumstances, a policy of discretion is likely to be less difficult to defend before the public than one based on commitment. Nonetheless, if a reasonably satisfactory way around these problems could be found, a policy based on commitment might be the more attractive. The aim of this paper is to suggest a simple and transparent way to derive an optimal monetary policy rule. Since the whole process is easily automated, it may provide a useful benchmark against which to compare other, probably far more labour intensive, policy choices. For recent related work on optimal monetary policy see Rudebusch (2001), Sack and Wieland (2000), Svensson (1997) and Taylor (1995, 2000).

Our proposal is to maximise an intertemporal quadratic objective function based on a trade-off between inflation and output subject to a dynamic constraint that is based on a vector autoregressive (VAR) model of the economy. The choice of a quadratic objective function in inflation and output (or the output gap) reflects common practice in the control and inflation targeting literatures, for example, Rudebusch and Svensson (1999) and Sack (2000). A more formal justification was provided by Rotemberg and Woodford (1999) - see also Woodford (2003) - who showed that such a quadratic function can be derived as an approximation to a micro founded macro model with standard preferences in terms of consumption.

More controversial is the choice of dynamic constraint relating the targeted variables to the policy instruments. Kydland and Prescott (1977), Barro and Gordon (1983) and Rogoff (1985) studied optimal monetary policy using rational expectations models of output and inflation. In contrast, Rudebusch and Svensson (1999) used a dynamic constraint obtained from an entirely backward looking model of output and inflation, while Clarida, Gali and Gertler (1999) and Rotemberg and Woodford (1998) employed a New Keynesian model. Optimal policy will clearly be affected by this choice. A more agnostic approach that seeks to avoid imposing a constraint not supported by the data is to use instead a data-based VAR. The logical drawback of using a VAR is that a change of policy rule will alter it. As a result, there must be a concern that any VAR is vulnerable to structural change. In principle, therefore, the Lucas Critique (Lucas, 1976) applies here. In practice, however, like Rudebusch (2002), we find that structural change to a VAR as a result of changing policy appears not to be much in evidence. Without wishing to claim that basing policy on VAR is a first-best approach compared with using a correctly-specified structural model, given the difficulty of agreeing on what that structural model should be, using a VAR may still provide a helpful benchmark against which to compare a first-best policy and any other policies such as one based on a Taylor rule or a policy of discretion. Moreover, using a VAR is consistent with Woodford's (2007) recommendation to base US monetary policy strategy on inflation forecast targeting whilst making no mention of how the forecasts should be obtained.

The idea of basing optimal monetary policy on a VAR is not new. Sack (2000) shows how this may be accomplished using a VAR in which the disturbances in the equations for the non-policy and policy variables are assumed to be uncorrelated. Martin and Salmon (1999) also use a VAR but with a different set of identifying restrictions from Sack. Having obtained the optimal policy rule, forecasts of the non-policy variables are derived from the VAR under control by replacing the original VAR equations for the policy instruments by the optimal policy rule. We refer to this methodology as the standard approach. Stock and Watson (2001) have made a related suggestion, namely to replace the interest rate equation in a VAR with a Taylor rule. This has the added drawback that the Taylor rule may not be an optimal choice.

Rather than make any assumptions about the correlation structure of the VAR disturbances, we estimate the VAR unrestrictedly and then derive the dynamic constraint relating the non-policy variables to the policy instruments by transforming the VAR so that the non-policy variables are conditioned on the policy instruments. Having derived the optimal policy rule, we construct the VAR under control by combining the sub-system of equations for the non-policy variables that make up the dynamic constraint with the optimal rule. We refer to this approach as the policy VAR (PVAR) method, where the PVAR is the resulting VAR under control.

Whereas in the optimal policy rule under the standard approach based on dynamic programming the policy instruments are related contemporaneously to the non-policy variables, in the PVAR approach the non-policy variables affect the policy instruments with a lag. For both methods the VAR under control has a lagged response to the policy instruments. Due to the structure of the optimal policy rule, in the standard method, the policy instruments are affected by the current disturbances to the non-policy VAR equations. If the identification is invalid due to correlation between the VAR disturbances of the non-policy and policy equations in the estimated VAR, then shocks to the non-policy variables in the VAR under control will be correlated with the disturbances of the original VAR equations for the policy instruments. This does not happen with the PVAR method.

There is, however, a potential problem with Sacks's approach. This is due to the particular identifying restrictions imposed on the VAR, notably, the assumption that the VAR disturbances of the equations for the non-policy and the policy variables are uncorrelated. It is common to introduce some minimal restrictions on a VAR in order to identify policy and other shocks. These usually take the form of assumed delays in the response of non-policy variables to policy shocks. A Choleski decomposition on the error covariance matrix of the VAR disturbances is the standard way to do this. Thus Sack assumes that there is a single-period lag in the response of non-policy variables to the policy instruments, and that the disturbances of the equations for the non-policy variables and the policy instruments are uncorrelated. Prominent studies along these lines are, for example, Sims (1980), Leeper et. al. (1996) and Christiano, Eichenbaum and Evans (1999). We refer to this as the standard approach. An alternative adopted by Bernanke and Blinder (1992) and Bernanke and Mihov (1998) is to start with a structural VAR which has uncorrelated disturbances, and then identify the resulting VAR by imposing zero restrictions on the coefficient matrix of the current variables.

As noted by Stock and Watson (2001), such assumptions are clearly more appropriate for high frequency macroeconomic data, such as daily or weekly than for lower frequency data, such as monthly, quarterly or annual data. The problem is that little macroeconomic data are available on a monthly basis, and often there is no good reason, beyond convenience, to assume that there are delays in the response to VAR shocks. For example, if the VAR is derived from a dynamic general equilibrium model, then a shock to households is most likely to affect simultaneously all of the households' decisions variables such as consumption, leisure and money demand. Moreover, financial markets and asset prices may be expected to response more or less instantaneously to policy shocks, see for example Bernanke and Mihov (1998) and Walsh (2003). For these reasons, but mainly because it is unnecessary to do so in order to formulate optimal policy, we prefer not to impose such identifying restrictions.

We argue that there are additional reasons for preferring the PVAR method. It is suitable for re-optimising policy each period based on up-dated estimates of the VAR as one would do in real-time applications and the optimal solutions for the state vector are the same whether one uses a state-space representation of the VAR under control, or a VAR based on the original variables. None of these holds for the standard approach as the outcomes for the non-policy variables next period, calculated from the VAR under control based on the orginal variables, are just the one-period ahead forecasts from the estimated VAR, and are different from the outcomes based on a state-space representation of the VAR under control. This first-period problem is akin to Woodford's (2003) timeless perspective.

We compare the two methods by deriving an optimal policy monetary rule using data for the US for the period 1960-2007. We examine alternative choices of the intertemporal welfare function based on flexible and strict inflation targeting, strict output targeting, interest rate smoothing and optimal flexible inflation targeting derived from Woodford's (2003) quadratic approximation to the type of intertemporal utility functions used in dynamic general equilibrium models. We also compare these methods with the use of a Taylor rule.

The paper is set out as follows. In section 2 we discuss how to formulate the dynamic constraint based on a VAR or a cointegrated VAR. In section 3 we derive the optimal policy rules for the standard and the PVAR methods and discuss some of the problems in implementing them. In section 4 we consider the problem of formulating optimal monetary policy for the US and the appropriate specification of the welfare function. Our results are reported in section 5 and our conclusions are stated in section 6. Our findings indicate that optimal monetary policy would have differed significantly from actual policy since 1991.

2 Formulating the dynamic constraint from a VAR

2.1 Stationary VAR

The aim of optimal policy is to respond to temporary and permanent shocks to the economy. The temporary shocks may be estimated from the disturbances of the original VAR. By evaluating the intertemporal welfare or cost function using the original data and with the forecasts derived from the VAR under control it is possible to compare the behaviour of the economy under control with not under control. Another way to compare the two is to examine the corresponding sets of dynamic multipliers. Permanent shocks are the due to changes in the non-stationary exogenous variables. This requires the use of a cointegrated VAR which we discuss below. Our discussion at this stage is based on all of the variables being stationary with the result that there are only temporary shocks.

The standard approach to formulating optimal policy based on a VAR used by Sack (2000) is to start by estimating a VAR which has equations for the target variables (those to be controlled), for the policy instruments (the control variables) and, optionally, for any other variables that might be involved in the transmission mechanism from the policy instruments to the policy targets. The second step is to derive the optimal rules by minimising an intertemporal cost function defined in terms of the target variables (and possibly also the policy instruments) subject to the constraint that the non-policy variables must satisfy their corresponding VAR equations. Thus the policy instruments are assumed to affect the non-policy variables with a lag. The welfare function is commonly chosen to be a quadratic function of the targets around their desired values. The optimal rules relate the instruments to the current values of the non-policy variables and the lagged state vector. The final step is to replace the original VAR equations for the policy instruments with optimal policy rule to form a new VAR in the original variables, the VAR under control.

This methodology is only valid if, as Sack assumes, the VAR disturbances of the non-policy and policy variables are uncorrelated. If the VAR disturbances are correlated, as will generally be the case, especially with quarterly data, a different method is required. This is because replacing the VAR equations for the policy variables with the control rule will alter the original correlation structure. Since the original disturbances of these equations are what optimal policy should be reacting to, this will not result in an optimal policy under control. Only if the VAR is restricted so that the disturbances of the non-policy equations are uncorrelated with those of the policy instrument equations will the standard procedure give the correct answer. If this is not a valid restriction then it would result in biased estimates of the shocks and an incorrect policy rule.

We suggest an alternative procedure when the disturbances are correlated which we call the policy VAR (PVAR) method, where the PVAR is the VAR under control. Rather than make any assumptions about the correlation structure of the VAR disturbances, we estimate the VAR unrestrictedly and then derive the dynamic constraint relating the non-policy variables to the policy instruments by transforming the VAR so that the non-policy variables are conditioned on the policy instruments. Having derived the optimal policy rule, we construct the VAR under control by combining the sub-system of equations for the non-policy variables that make up the dynamic constraint with the optimal rule. The disturbances in the non-policy equations are now different from those in the original VAR and not the same as in the standard method. In the PVAR these disturbances are the projection of the original disturbances onto the space orthogonal to the disturbances of the policy equation disturbances.

We examine the details more closely. Consider the VAR(p)

$$\mathbf{z}_t = \mathbf{a} + \mathbf{A}(L)\,\mathbf{z}_{t-1} + \mathbf{e}_t, \tag{1}$$

where $\mathbf{z}_t' = (\mathbf{z}_{1,t}', \mathbf{z}_{2,t}')$, $\mathbf{z}_{1,t}$ is an $s \times 1$ vector of non-policy or state variables, $\mathbf{z}_{2,t}$ is a $c \times 1$ vector of policy instruments or control variables, $q = s + c$, $\mathbf{a}$ is a vector of constant terms, $\mathbf{A}(L) = \Sigma_{i=0}^{p-1}\mathbf{A}_i L^i$, with $\mathbf{A}_i$ indicating a $q \times q$ matrix of lag i coefficients, L is the lag operator, $\mathbf{e}_t$ is a vector of stochastic disturbances with $E[\mathbf{e}_t] = \mathbf{0}$, $E[\mathbf{e}_t\mathbf{e}_t'] = \mathbf{\Sigma}$ and $E[\mathbf{e}_t\mathbf{e}_{t-i}'] = 0$ for $i > 0$. We assume that $\mathbf{\Sigma}$ is unrestricted, thereby allowing the VAR disturbances to be correlated. The VAR can be partitioned as

$$\begin{bmatrix} \mathbf{z}_{1,t} \\ \mathbf{z}_{2,t} \end{bmatrix} = \begin{bmatrix} \mathbf{a}_{10} \\ \mathbf{a}_{20} \end{bmatrix} + \begin{bmatrix} \mathbf{A}_{11}(L) & \mathbf{A}_{12}(L) \\ \mathbf{A}_{21}(L) & \mathbf{A}_{22}(L) \end{bmatrix} \begin{bmatrix} \mathbf{z}_{1,t-1} \\ \mathbf{z}_{2,t-1} \end{bmatrix} + \begin{bmatrix} \mathbf{e}_{1t} \\ \mathbf{e}_{2t} \end{bmatrix} \tag{2}$$

where $\mathbf{A}_{11}(L) = \Sigma_{i=0}^{p-1}\mathbf{A}_{11i}L^i$, $\mathbf{A}_{12}(L) = \Sigma_{i=0}^{p-1}\mathbf{A}_{12i}L^i$, $\mathbf{A}_{21}(L) = \Sigma_{i=0}^{p-1}\mathbf{A}_{21i}L^i$ and $\mathbf{A}_{22}(L) = \Sigma_{i=0}^{p-1}\mathbf{A}_{22i}L^i$. $\mathbf{e}_t' = (\mathbf{e}_{1,t}', \mathbf{e}_{2,t}')$ and the covariance matrix of the disturbances is partitioned conformably as

$$\mathbf{\Sigma} = \begin{bmatrix} \mathbf{\Sigma}_{11} & \mathbf{\Sigma}_{12} \\ \mathbf{\Sigma}_{21} & \mathbf{\Sigma}_{22} \end{bmatrix}.$$

where $\mathbf{\Sigma}_{12} = \mathbf{\Sigma}_{21}'$. Sack (2000) imposes the restriction $\mathbf{\Sigma}_{12} = 0$, implying that shocks to his single policy instrument (including unanticipated policy changes) have no contemporaneous effect on the non-policy variables.

Bernanke and Blinder (1992), whose purpose was to identify the shocks rather than perform optimal control analysis, started not with equation (1), but with a structural VAR, i.e. with the SVAR

$$\mathbf{B}\mathbf{z}_t = \mathbf{a} + \mathbf{A}\left(L\right)\mathbf{z}_{t-1} + \mathbf{u}_t,$$

where the disturbances $\mathbf{u}_t$ of the policy and non-policy instruments are assumed to be uncorrelated. They then consider two possible identification schemes: partitioning $\mathbf{B}$, they set either $\mathbf{B}_{12}$ or $\mathbf{B}_{21}$ equal to zero. If the variables in z_t are ordered non-policy and policy as before then this implies, respectively, that either the policy variables affect the non-policy variables with a lag or vice-versa. Bernanke and Mihov (1998) argue in favour of the restriction $\mathbf{B}_{12}= \mathbf{0}$. They also point out that even this restriction may not be suitable if the data period is so long that the non-policy variables have time to react to the policy instruments within the period of observation. The restrictions imposed by Sack, Bernanke and Blinder, and Bernanke and Mihov are sometimes known as partial identification of a VAR, or using a semi-structural VAR. Martin and Salmon (1999), who also consider optimal policy with a VAR, argue that identifying a VAR through the sort of recursive restrictions used by Sack is unsatisfactory. Instead they use selective contemporaneous non-recursive restrictions to the disturbances.

In contrast, we impose no identification restrictions on the state-vector equations of the VAR and so allow $\mathbf{\Sigma}_{12} \neq 0$, implying that shocks to the policy instrument (including unanticipated policy changes) may affect the non-policy variables contemporaneously. If $\mathbf{e}_{1,t}$ is correlated with $\mathbf{e}_{2,t}$ then we can express their relation through

$$\mathbf{e}_{1,t} = \boldsymbol{\epsilon}_t + \mathbf{G}\mathbf{e}_{2,t},$$

where $\boldsymbol{\epsilon}_t$ is the component of $\mathbf{e}_{1,t}$ that is uncorrelated with $\mathbf{e}_{2,t}$. Consequently, we may write

$$\mathbf{e}_t = \begin{bmatrix} \mathbf{e}_{1,t} \\ \mathbf{e}_{2,t} \end{bmatrix} = \begin{bmatrix} \mathbf{I} & \mathbf{G} \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \boldsymbol{\epsilon}_{1t} \\ \mathbf{e}_{2,t} \end{bmatrix} = \mathbf{H} \begin{bmatrix} \boldsymbol{\epsilon}_t \\ \mathbf{e}_{2,t} \end{bmatrix}.$$

It follows from $E[\mathbf{e}_t\mathbf{e}_t'] = \boldsymbol{\Sigma}$ that $\boldsymbol{\Sigma}_{12} = \mathbf{G}\boldsymbol{\Sigma}_{22}$, implying that $\mathbf{G} = \boldsymbol{\Sigma}_{12}\boldsymbol{\Sigma}_{22}^{-1}$.

We now pre-multiply the VAR by the transformation matrix

$$\mathbf{H}^{-1} = \begin{bmatrix} \mathbf{I} & -\mathbf{G} \\ \mathbf{0} & \mathbf{I} \end{bmatrix}.$$

to obtain

$$\mathbf{H}^{-1}\mathbf{z}_t = \mathbf{H}^{-1}\mathbf{a} + \mathbf{H}^{-1}\mathbf{A}\,(L)\,\mathbf{z}_{t-1} + \mathbf{H}^{-1}\mathbf{e}_t,$$

This defines two sets of equations with uncorrelated disturbances:

$$\begin{aligned} \mathbf{z}_{1,t} &= [\mathbf{a}_{10} - \mathbf{G}\mathbf{a}_{20}] + \mathbf{G}\mathbf{z}_{2,t} + [\mathbf{A}_{11}\,(L) - \mathbf{G}\mathbf{A}_{21}\,(L)]\,\mathbf{z}_{1,t-1} + \\ &\quad + [\mathbf{A}_{12}\,(L) - \mathbf{G}\mathbf{A}_{22}\,(L)]\,\mathbf{z}_{2,t-1} + \boldsymbol{\epsilon}_t \qquad (3) \\ \mathbf{z}_{2,t} &= \mathbf{a}_{20} + \mathbf{A}_{21}\,(L)\,\mathbf{z}_{1,t-1} + \mathbf{A}_{22}\,(L)\,\mathbf{z}_{2,t-1} + \mathbf{e}_{2,t}. \qquad (4) \end{aligned}$$

Hence, from equation (3), under the PVAR method the state vector $\mathbf{z}_{1,t}$ is conditioned on the control vector $\mathbf{z}_{2,t}$ which is uncorrelated with $\boldsymbol{\epsilon}_{1t}$. Equation (3) describes the law of motion of the state vector of non-policy variables, conditional on the policy instruments. The optimal policy is derived subject to this equation, not the equation for $\mathbf{z}_{1,t}$ in the original VAR, equation (2).

2.2 Cointegrated VAR

If the data are non-stationary, or a mixture of stationary and non-stationary variables, then we have a choice. We could just ignore the presence of non-stationary variables and still use a conventional levels VAR to represent the data. Or we could estimate a cointegrated VAR and then re-write this as a levels VAR. And if there is no cointegration we could first difference the variables and then estimate a levels VAR.

If we ignore the presence of cointegration we would, of course, be throwing away information which could be used to improve the efficiency of the coefficient estimates of the VAR when re-written in levels. In particular, we would be ignoring information on how many cointegrating vectors there are. This is information that could be incorporated in the estimation method as in the Johansen estimator. Further, we know that where there is more than one cointegrating vector they cannot be given an economic interpretation even though they are identified in a statistical sense, see Wickens (1996). Wickens and Motto (2003) have suggested a way of identifying the cointegrating vectors that imposes minimal restrictions. This consists of using long-run structural information to restrict the cointegrating vectors and, optionally, classifying the variables as either endogenous or exogenous.

Henceforth, however, we will proceed as though the variables are stationary.

3 Optimal policy with a VAR

3.1 The problem

The optimal control of a time-separable inter-temporal quadratic objective function constrained by a stochastic linear dynamic system is well known. The solution may be obtained either by using the method of dynamic programming or the method of Lagrange multipliers; both techniques lead to the same solution.[1] When dynamic programming is used the problem is commonly referred to as linear quadratic dynamic programming. We wish to compare the optimal solution based on the PVAR method with that based on the standard method using a VAR. We therefore derive the solutions for the policy variables $\mathbf{z}_{2t}$ for the case where the dynamic constraint determining the non-policy variables $\mathbf{z}_{1t}$ is the conditional VAR, equation (3), and for the case where the equations for $\mathbf{z}_{1t}$ are those in the VAR, equation (1).

[1] An accurate comparison of the use of dynamic programming and Lagrange multipliers tecniques for the assessment of optimal policy rules can be found in Chow (1976).

We assume that the quadratic loss function of the policy maker is

$$L_t = E_t \sum_{s=0}^{\infty} \beta^s \left[(\mathbf{y}_{t+s} - \bar{\mathbf{y}})' \mathbf{W} (\mathbf{y}_{t+s} - \bar{\mathbf{y}}) \right], \tag{5}$$

where $\mathbf{y}_t' = \begin{bmatrix} \mathbf{z}_{1t}' & \ldots & \mathbf{z}_{1,t-p}' & \mathbf{z}_{2t}' & \ldots & \mathbf{z}_{2,t-p}' \end{bmatrix}$, $\bar{\mathbf{y}}$ is a target vector and $\mathbf{W}$ is a symmetric positive semi-definite matrix of policy weights which are assumed to be constant. This formulation is sufficiently general to allow the objective function to include first differences of the variables. This will only affect the form of the matrix $\mathbf{W}$.

The value function $V(\mathbf{y}_t)$, i.e. the minimum value at time t of the welfare loss under the infinite sequence of controls $\{\mathbf{z}_{2t+s}\}_{s=0}^{\infty}$, is given by

$$V(\mathbf{y}_t) = \min_{\{\mathbf{z}_{2t+s}\}_{s=0}^{\infty}} E_t \sum_{s=0}^{\infty} \beta^s \left[(\mathbf{y}_{t+s} - \bar{\mathbf{y}})' \mathbf{W} (\mathbf{y}_{t+s} - \bar{\mathbf{y}}) \right]$$

Using the Bellman (1957) principle, the value function can be re-written in the recursive form

$$V(\mathbf{y}_t) = \min_{\{\mathbf{z}_{2t+s}\}_{s=0}^{\infty}} (\mathbf{y}_t - \bar{\mathbf{y}})' \mathbf{W} (\mathbf{y}_t - \bar{\mathbf{y}}) + \beta E_t \left[V(\mathbf{y}_{t+1}) \right] \tag{6}$$

3.2 Standard VAR

The dynamic constraint in the standard approach is based on the sub-system of equations for $\mathbf{z}_{1t}$ in the original VAR, equation (1). In state-space (companion) form it can be re-written as

$$\mathbf{y}_t = \widetilde{\mathbf{a}} + \widetilde{\mathbf{A}} \mathbf{y}_{t-1} + \widetilde{\mathbf{B}} \mathbf{z}_{2,t-1} + \mathbf{v}_t \tag{7}$$

where $E_t \left[\mathbf{z}_{2,t-1}' \mathbf{v}_t \right] = \mathbf{0}$. Maximising the value function (6) subject to (7) gives the optimal rule, see Sack(2000) and Ljungqvist and Sargent (2004). The optimal solution is

$$\begin{aligned} \mathbf{z}_{2t} &= \tilde{\mathbf{f}} + \tilde{\mathbf{F}} \mathbf{y}_t, \\ \tilde{\mathbf{F}} &= -\left(\widetilde{\mathbf{B}}' \tilde{\mathbf{P}} \widetilde{\mathbf{B}} \right)^{-1} \widetilde{\mathbf{B}}' \tilde{\mathbf{P}} \widetilde{\mathbf{A}} \\ \tilde{\mathbf{f}} &= -\left(\widetilde{\mathbf{B}}' \tilde{\mathbf{P}} \widetilde{\mathbf{B}} \right)^{-1} \widetilde{\mathbf{B}}' \left(\tilde{\mathbf{p}} - \tilde{\mathbf{P}} \widetilde{\mathbf{a}} \right), \end{aligned} \tag{8}$$

where $\tilde{\mathbf{P}}$ and $\tilde{\mathbf{p}}$ satisfy the time-invariant Riccati equations

$$\tilde{\mathbf{P}} = \mathbf{W}+\beta\widetilde{\mathbf{A}}'\tilde{\mathbf{P}}\widetilde{\mathbf{A}}-\beta\widetilde{\mathbf{A}}'\tilde{\mathbf{P}}\widetilde{\mathbf{B}}\left(\widetilde{\mathbf{B}}'\tilde{\mathbf{P}}\widetilde{\mathbf{B}}\right)^{-1}\widetilde{\mathbf{B}}'\tilde{\mathbf{P}}\widetilde{\mathbf{A}} \tag{9}$$

$$\tilde{\mathbf{p}} = \mathbf{W}\bar{\mathbf{y}}+\beta\left(\widetilde{\mathbf{A}}'+\tilde{\mathbf{F}}'\widetilde{\mathbf{B}}'\right)\left(\tilde{\mathbf{p}}-\tilde{\mathbf{P}}\widetilde{\mathbf{a}}\right) \tag{10}$$

Hence, in the standard method, equation (8), the policy instrument responds contemporaneously to $\mathbf{y}_t$. Note that this solution differs from that of Chow (1976), pp. 156-160 and 176-178 due to the presence of the discount factor β. The loss function, equation (5), differs from Chow's which replaces $\beta^s\mathbf{W}$ by the more general $\mathbf{W}_s$. Equation (9) is non-linear but satisfies a fixed-point theorem, the solution for $\mathbf{P}$ must be therefore be obtained through numerical iteration.

The maximised value function is now given by

$$V(\mathbf{y}_t) = \mathbf{y}_t'\tilde{\mathbf{P}}\mathbf{y}_t + 2\mathbf{y}_t'\tilde{\mathbf{p}} + \widetilde{\mathbf{c}} \tag{11}$$

Substituting (8) into (7) gives the state-space representation

$$\mathbf{y}_t = \widetilde{\mathbf{r}} + \widetilde{\mathbf{R}}\mathbf{y}_{t-1} + \mathbf{v}_t \tag{12}$$

where $\widetilde{\mathbf{r}}_t = \widetilde{\mathbf{a}} + \widetilde{\mathbf{B}}\widetilde{\mathbf{f}}$ and $\widetilde{\mathbf{R}} = \widetilde{\mathbf{A}} + \widetilde{\mathbf{B}}\widetilde{\mathbf{F}}$.

Although commonly used to calculate the outcomes for the non-policy variables, equation (12) must be interpreted with care. It determines the behaviour of the state vector on the implicit assumption that the policy instruments have always been generated by the above policy rule, and there has been no switch of policy. This has echoes of Woodford's timeless perspective in which the different behaviour of the non-policy variables in first period is ignored, see Woodford (2003). If the optimal policy is implemented at time t then past values of the state vector, which are required to calculate $\mathbf{y}_t$ in equation (24), will be determined prior to the change of policy and not, as implicitly assumed, by equation (24).

Further appreciation of this problem may be acquired by considering an alternative way of writing the VAR under control. Written in terms of the original variables, the control rule in the standard approach takes the form

$$\mathbf{\Lambda z}_{1t}+\mathbf{z}_{2t}=\widetilde{\mathbf{a}}_{20}+\widetilde{\mathbf{A}}_{21}(\mathbf{L})\mathbf{z}_{1,t-1}+\widetilde{\mathbf{A}}_{22}(\mathbf{L})\mathbf{z}_{2,t-1} \tag{13}$$

Thus it is necessary for the policy maker to know the non-policy variables z_{1t} in order to set policy. The VAR under control for the standard approach is obtained by combining the original VAR equations for $\mathbf{z}_{1t}$ with the policy rule to give

$$\begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{\Lambda} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{z}_{1t} \\ \mathbf{z}_{2t} \end{bmatrix} = \begin{bmatrix} \mathbf{a}_{10} \\ \widetilde{\mathbf{a}}_{20} \end{bmatrix} + \begin{bmatrix} \mathbf{A}_{11}(L) & \mathbf{A}_{12}(L) \\ \widetilde{\mathbf{A}}_{21}(L) & \widetilde{\mathbf{A}}_{22}(L) \end{bmatrix} \begin{bmatrix} \mathbf{z}_{1,t-1} \\ \mathbf{z}_{2,t-1} \end{bmatrix} + \begin{bmatrix} \mathbf{e}_{1t} \\ 0 \end{bmatrix} \tag{14}$$

This can be re-written as the VAR under control by pre-multiplying by the inverse of the matrix of coefficients of $\mathbf{z}_t$,

$$\begin{bmatrix} \mathbf{z}_{1t} \\ \mathbf{z}_{2t} \end{bmatrix} = \begin{bmatrix} \mathbf{a}_{10} \\ -\mathbf{\Lambda a}_{10}+\widetilde{\mathbf{a}}_{20} \end{bmatrix} + \begin{bmatrix} \mathbf{A}_{11}(L) & \mathbf{A}_{12}(L) \\ -\mathbf{\Lambda A}_{11}(L)+\widetilde{\mathbf{A}}_{21}(L) & -\mathbf{\Lambda A}_{12}(L)+\widetilde{\mathbf{A}}_{22}(L) \end{bmatrix} \begin{bmatrix} \mathbf{z}_{1,t-1} \\ \mathbf{z}_{2,t-1} \end{bmatrix} + \begin{bmatrix} \mathbf{e}_{1t} \\ -\mathbf{\Lambda e}_{1t} \end{bmatrix} \tag{15}$$

It can be shown that the equation for the non-policy variables z_{1t} in equation (15) is different from that in equation (12). This is entirely due to the presence of the current value of $\mathbf{z}_{1t}$ in the control rule, equation (13) which requires that z_{1t} must be known in order to set policy, and then z_{2t} affects $z_{1,t+1}$. Hence, equation (15), and not equation (12), should be used following a switch of policy.

The expected loss under the standard approach may be evaluated as before. It is

$$L_t = \sum_{s=0}^{\infty} \beta^s \left[(E_t\mathbf{y}_{t+s}-\bar{\mathbf{y}})' \mathbf{W} (E_t\mathbf{y}_{t+s}-\bar{\mathbf{y}}) \right] + \sum_{s=0}^{\infty} \beta^s tr\mathbf{W}\widetilde{\mathbf{\Gamma}}_s \tag{16}$$

where

$$
\begin{aligned}
\widetilde{\mathbf{\Gamma}}_s &= E_t(\mathbf{y}_{t+s} - E_t\mathbf{y}_{t+s})(\mathbf{y}_{t+s} - E_t\mathbf{y}_{t+s})' \\
&= (\mathbf{I} + \widetilde{\mathbf{R}} + ... + \widetilde{\mathbf{R}}^{s-1})\widetilde{\mathbf{\Omega}}(\mathbf{I} + \widetilde{\mathbf{R}} + ... + \widetilde{\mathbf{R}}^{\mathbf{s-1}})' \\
&= \widetilde{\mathbf{R}}\widetilde{\mathbf{\Gamma}}_{s-1}\widetilde{\mathbf{R}}' + \widetilde{\mathbf{\Omega}}
\end{aligned}
$$

$\widetilde{\mathbf{\Gamma}}_0 = \mathbf{0}$ and $E\mathbf{v}_t\mathbf{v}_t' = \widetilde{\mathbf{\Omega}}$. Thus $\mathbf{\Gamma}_s$ is the conditional covariance matrix of $\mathbf{y}_{t+s}$ given information at time t. Equation (16) shows that the expected welfare cost can be decomposed into two parts. The first term on the right hand side of equation (16) is the deterministic component of the welfare cost and measures the cost of having a long-run target different from the conditional expectation of the vector $\mathbf{y}_{t+s}$. The second term is the stochastic component of the welfare cost, which depends upon the volatility of the vector $\mathbf{y}_{t+s}$. In particular, $\mathbf{\Gamma}_s$ measures the volatility of the forecast error due to the presence of random disturbances $\mathbf{e}_{1,t+s}$ which cause deviations of $\mathbf{y}_{t+s}$ from its expected path. $\mathbf{\Gamma}_s$ may be obtained from equation (12).

3.3 PVAR method

In the PVAR approach the dynamic constraint is based on equation (3). In state-space form it can be re-written as

$$
\mathbf{y}_t = \widehat{\mathbf{a}} + \widehat{\mathbf{A}}\mathbf{y}_{t-1} + \widehat{\mathbf{B}}\mathbf{z}_{2t} + \mathbf{u}_t \tag{17}
$$

where $E_t\left[\mathbf{z}_{2t}'\mathbf{u}_t\right] = \mathbf{0}$, see the appendix for further details. It can be shown that maximising (6) subject to (17) gives the optimal rule

$$
\begin{aligned}
\mathbf{z}_{2t} &= \mathbf{f} + \mathbf{F}\mathbf{y}_{t-1}, \qquad (18) \\
\mathbf{F} &= -\left(\widehat{\mathbf{B}}'\mathbf{P}\widehat{\mathbf{B}}\right)^{-1}\widehat{\mathbf{B}}'\mathbf{P}\widehat{\mathbf{A}} \qquad (19) \\
\mathbf{f} &= -\left(\widehat{\mathbf{B}}'\mathbf{P}\widehat{\mathbf{B}}\right)^{-1}\widehat{\mathbf{B}}'\left(\mathbf{P}\widehat{\mathbf{a}} - \mathbf{p}\right) \qquad (20)
\end{aligned}
$$

where $\mathbf{P}$ and $\mathbf{p}$ satisfy the time-invariant Riccati equations

$$\mathbf{P} = \mathbf{W}+\beta\widehat{\mathbf{A}}'\mathbf{P}\widehat{\mathbf{A}}-\beta\widehat{\mathbf{A}}'\mathbf{P}\widehat{\mathbf{B}}\left(\widehat{\mathbf{B}}'\mathbf{P}\widehat{\mathbf{B}}\right)^{-1}\widehat{\mathbf{B}}'\mathbf{P}\widehat{\mathbf{A}} \tag{21}$$

$$\mathbf{p} = \mathbf{W}\bar{\mathbf{y}}+\beta\left(\widehat{\mathbf{A}}'+\mathbf{F}'\widehat{\mathbf{B}}'\right)(\mathbf{p}-\mathbf{P}\widehat{\mathbf{a}}) \tag{22}$$

The maximised value function is then given by

$$V(\mathbf{y}_t) = \mathbf{y}_t'\mathbf{P}\mathbf{y}_t + 2\mathbf{y}_t'\mathbf{p} + \mathbf{c} \tag{23}$$

The behaviour of the state vector under control is usually expressed in state-space form. It is obtained by substituting the optimal rule (18) into equation (17) to give

$$\mathbf{y}_t = \mathbf{r} + \mathbf{R}\mathbf{y}_{t-1} + \mathbf{u}_t, \tag{24}$$

where $\mathbf{r} = \widehat{\mathbf{a}} + \widehat{\mathbf{B}}\mathbf{f}$ and $\mathbf{R} = \widehat{\mathbf{A}} + \widehat{\mathbf{B}}\mathbf{F}$. Although equation (24) has the same form as equation (12), it does not have the same drawback. To see this we express the VAR under control in terms of the original variables. The control rule is

$$\mathbf{z}_{2t} = \mathbf{a}_{20}^{*} + \mathbf{A}_{21}^{*}(L)\mathbf{z}_{1,t-1} + \mathbf{A}_{22}^{*}(L)\mathbf{z}_{2t-1}$$

where the policy instrument is based on information prior to period t. The VAR under control -

the policy VAR or PVAR - may be written as

$$\begin{aligned}
\begin{bmatrix} \mathbf{z}_{1t} \\ \mathbf{z}_{2t} \end{bmatrix} &= \begin{bmatrix} \mathbf{I} & \mathbf{G} \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{a}_{10} - \mathbf{G}\mathbf{a}_{20} \\ \mathbf{a}_{20}^{*} \end{bmatrix} + \\
&+ \begin{bmatrix} \mathbf{I} & \mathbf{G} \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{A}_{11}(L) - \mathbf{G}\mathbf{A}_{21}(L) & \mathbf{A}_{12}(L) - \mathbf{G}\mathbf{A}_{22}(L) \\ \mathbf{A}_{21}^{*}(L) & \mathbf{A}_{22}^{*}(L) \end{bmatrix} \begin{bmatrix} \mathbf{z}_{1,t-1} \\ \mathbf{z}_{2,t-1} \end{bmatrix} \\
&+ \begin{bmatrix} \mathbf{I} & \mathbf{G} \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \boldsymbol{\epsilon}_t \\ 0 \end{bmatrix} \\
&= \begin{bmatrix} \mathbf{a}_{10} - \mathbf{G}(\mathbf{a}_{20} - \mathbf{a}_{20}^{*}) \\ \mathbf{a}_{20}^{*} \end{bmatrix} + \\
&\quad \begin{bmatrix} \mathbf{A}_{11}(L) - \mathbf{G}[\mathbf{A}_{21}(L) - \mathbf{A}_{21}^{*}(L)] & \mathbf{A}_{12}(L) - \mathbf{G}[\mathbf{A}_{22}(L) - \mathbf{A}_{22}^{*}(L)] \\ \mathbf{A}_{21}^{*}(L) & \mathbf{A}_{22}^{*}(L) \end{bmatrix} \begin{bmatrix} \mathbf{z}_{1,t-1} \\ \mathbf{z}_{2,t-1} \end{bmatrix} \\
&+ \begin{bmatrix} \boldsymbol{\epsilon}_t \\ 0 \end{bmatrix} \qquad (25)
\end{aligned}$$

In contrast to the standard approach, where the equation for the non-policy variables z_{1t} in equation (15) is different from that in equation (12), it can be shown that the equation for z_{1t} in equation (25) is identical to that in equation (24). In other words, in calculating the outcomes for the variables under control, the PVAR approach does not suffer from the ambiguities of the standard approach.

On the other hand, the new equation for the non-policy variables z_{1t} in equation (25) differs from the original VAR if the new policy rule differs from the original VAR equations for the policy variables. Equation (25) shows that there is a lag in the effect of the policy instrument on the non-policy variables. Hence, in effect, policy is based on the expected future values of the non-policy variables next period. It therefore accords with Woodford's (2007) recommendation to base monetary policy on forecasting inflation. We also note that as a result of switching to the optimal policy the coefficients of the VAR have changed. This shows that the Lucas Critique applies even when the original model is a VAR. Further, we note that the VAR under control has a singular error covariance matrix because the new VAR equations for the policy variables are deterministic.

The VAR under control in the PVAR approach differs from that for the standard approach in two further ways. First, in the standard approach, because policy requires a knowledge of the current non-policy variables z_{1t}, the equations for the policy instruments in the VAR under control have a disturbance term $\mathbf{e}_{1t}$ that is perfectly correlated with disturbances in the policy equations. In the PVAR the policy equations have no disturbance term. Since there is a one period lag in the response of the non-policy variables and future disturbances would be unknown in practice, setting these disturbances to zero implies that, as in the PVAR method, in effect, policy is once again based on the expected future values of the non-policy variables next period.

The second difference is that, in the standard approach the dynamic structure of the equations for the non-policy variables are unaffected when under control, whereas they are altered in the PVAR. This implies that re-optimising each period and building back the disturbances $\mathbf{e}_{1t}$ would result in no difference in the non-policy variables from their actual values even though the values of the policy instruments would be different from the actual data. This point is pertinent for our empirical experiments below.

The loss function, equation (5), may be evaluated under control by re-writing it as

$$L_t = \sum_{s=0}^{\infty} \beta^s \left[(E_t \mathbf{y}_{t+s} - \bar{\mathbf{y}})' \mathbf{W} (E_t \mathbf{y}_{t+s} - \bar{\mathbf{y}}) \right] + \sum_{s=0}^{\infty} \beta^s tr \mathbf{W} \mathbf{\Gamma}_s \tag{26}$$

where $\mathbf{\Gamma}_s$ now measures the volatility of the forecast error due to the presence of random disturbances $\boldsymbol{\epsilon}_{t+s}$ which cause deviations of $\mathbf{y}_{t+s}$ from its expected path. $\mathbf{\Gamma}_s$ may be obtained from equation (1). Denoting $E\mathbf{u}_t\mathbf{u}_t' = \Omega$, then, as $\mathbf{u}_{t+s}$ has a constant variance,

$$\begin{aligned} \mathbf{\Gamma}_s &= E_t(\mathbf{y}_{t+s} - E_t\mathbf{y}_{t+s})(\mathbf{y}_{t+s} - E_t\mathbf{y}_{t+s})' \\ &= (\mathbf{I} + \mathbf{R} + ... + \mathbf{R}^{s-1})\mathbf{\Omega}(\mathbf{I} + \mathbf{R} + ... + \mathbf{R}^{\mathbf{s-1}})' \\ &= \mathbf{R}\mathbf{\Gamma}_{s-1}\mathbf{R}' + \Omega \end{aligned}$$

where $\mathbf{\Gamma}_0 = \mathbf{0}$.

To summarise, if the original VAR disturbances are correlated, then the PVAR method should be followed instead of the standard approach which assumes that they are uncorrelated, and hence that $\mathbf{G} = \mathbf{0}$. The optimal rules are then different, as are the behaviour of the non-policy variables under control and the expected losses. The disturbances of a VAR under control have been shown to be different in the two methods. Using the standard method the covariance matrix of the disturbances of the non-policy and the policy equations will be perfectly correlated; under the PVAR method the covariance matrix of these disturbances will be singular. We have also shown that for the standard approach the calculation of the outcomes under control using the state-space VAR and the VAR under control written in terms of the original variables differ, whereas for the PVAR the two methods give the same solution. This implies that for the standard approach, in order to deal with the initial effects of a switch to the optimal policy, the VAR under control should be used rather than the state-space VAR. However, if, for the standard approach, policy is re-optimised each period, and so only the first period is implemented, then using the VAR under control rather than the state-space representation, implies that the outcomes for the non-policy variables are unaffected by policy as they are simply the one-period ahead forecasts from the original VAR. All of this makes the case for preferring the PVAR approach compelling.

3.4 Options for implementing optimal policy

Our policy experiments all assume that policy is re-optimised each period. This implies that we only need to consider the first period following the execution of the new policy. There are several ways to implement this which differ in the amount of information used.

Option1

This option makes the artificial assumption that the VAR coefficients are known for the whole sample period, even though the VAR cannot be estimated until after this period has passed. It is therefore a counter-factual experiment. It is also the experiment considered by Sack. Thus, we estimate the VAR for the whole data period and then re-construct the VAR under control for both the standard and the PVAR policy rules by computing the optimal values of the policy instruments and the one-period ahead forecasts of the policy target variables using actual past values of the state vector $\mathbf{z}_{t-1}$ and ignoring the period t disturbances $\mathbf{e}_{1t}$ (ST) and $\boldsymbol{\epsilon}_t$ (PVAR). This gives expected one-period-ahead forecasts. For the standard approach these forecasts are obtained from equation (15) and hence are just the one-period ahead in-sample forecasts from the original VAR; the switch of policy has therefore had no effect. If the state-space representation, equation (12) were used instead, the one-period ahead forecasts under control would differ from the in-sample forecasts from the original VAR. In contrast, the PVAR forecasts under control will differ from in-sample forecasts from the original VAR.

Option 2

This is the same as Option 1 except the period t disturbances $\mathbf{e}_{1t}$ (ST) and $\boldsymbol{\epsilon}_t$ (PVAR) are added to the VAR forecasts. This gives the "actual" outcomes under control. For the standard approach the outcomes under control are identical to the actual values because actual data is used for past values of the non-policy variables in the VAR under control. For the PVAR the outcome for the policy instrument (the Federal Funds rate) is the same as for option 1, but the equations for the non-policy variables in the VAR under control are altered.

Option 3

A step towards great realism is to estimate the VAR recursively adding one period at a time before re-constructing the VAR under control for both the standard and the PVAR policy rules. For this option, where policy is re-optimised each period, it is even more appropriate to use the VAR under control and not the state-space representation under control. In this option the data used at each recursion are the actual values observed and not the variables under control. This corresponds to Options 1 and 2 where the actual data are used to estimate the VAR. The rest is the same as in Option 1. In particular, we obtain expected one-period-ahead forecasts. For the standard approach the one-period ahead forecasts are the same as the those from the recursively estimated VAR and hence are also unaffected by a switch of policy.

Option 4

This is the same as Option 3 except that we include the period t disturbances $\mathbf{e}_{1t}$ (ST) and $\boldsymbol{\epsilon}_t$ (PVAR) to give the "actual" outcomes under control. For the standard approach the outcomes under control are again identical to the actual values because actual data is used for past values of the non-policy variables in the VAR under control and, in addition, actual data are being used to estimate the VAR. For the PVAR the outcome for the policy instrument is the same as for option 3, but the equations for the non-policy variables in the VAR under control are different.

It is clear that using a VAR under control rather than the state-space representation under control and assuming re-optimisation each period renders the standard approach useless. For the standard approach to possess any value it is necessary to use the state-space representation under control. This is what Sack does. Hence, given the logical problems associated with using the state-space representation, only the PVAR approach has any practical value. This argument is reinforced if we reflect that in practice - i.e. in real time - we would re-estimate the VAR each period and then re-optimise each period.

4 Optimal monetary policy

We now apply the PVAR method to the problem of optimal monetary policy based on using the official interest rate as the monetary policy instrument where the monetary authority is either a strict or a flexible inflation targeter. We also allow the monetary authority to smooth the official interest rate, if it wishes, in order to avoid instrument instability, see Goodhart (1998), Rudebusch (1998) and Sack (2000). We therefore consider the following quadratic intertemporal cost function:

$$L = E_t \sum_{s=0}^{\infty} \beta^{t+s} \left[\lambda_\pi \left(\pi_{t+s} - \pi_t^* \right)^2 + \lambda_y (y_{t+s} - y_t^*)^2 + \lambda_{\Delta rs} \Delta rs_{t+s}^2 \right], \tag{27}$$

where π_t is the rate of inflation, y_t is the output gap, rs_t is a nominal short-term interest rate, and is the official interest rate controlled by the monetary authority. π_t^* and y_t^* are target levels. We set $\lambda_\pi = 1$ throughout. If in addition $\lambda_y = \lambda_{\Delta rs} = 0$ then the monetary authority is a strict inflation targeter. Setting $\lambda_y > 0$ implies that the monetary authority is a flexible inflation targeter and setting $\lambda_{\Delta rs} > 0$ implies that the monetary authority attaches a cost to changing the official interest rate and so prefers to smooth interest rates. We consider the effects of different choices of λ_y and $\lambda_{\Delta rs}$.[2]

In order to evaluate the effect on optimal policy of different choices of the weights λ_y and $\lambda_{\Delta rs}$, like Rudebusch and Svensson (1999), we consider a range of values. These are summarised in table 1. The choices range from strict inflation targeting where $\lambda_y = \lambda_{\Delta rs} = 0$, to flexible targeting where $\lambda_y \neq 0$, to interest rate smoothing where $\lambda_{\Delta rs} \neq 0$. An alternative approach would be to calibrate the quadratic approximation to a standard intertemporal utility as derived by Woodford (2003); see also Wickens (2008). Starting with the instantaneous utility function

$$U_t = \ln c_t - \gamma \ln y_t(z) \tag{28}$$

where c_t is an index of total consumption given by the CES function

$$c_t = \left[\int_0^1 c_t(z)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}, \quad \sigma > 1$$

and the last term in the utility function reflects utility from leisure which is inversely related to the work effort required to produce $y_t(z)$ of good z. If, under Calvo pricing, a proportion ρ of firms can change their price in any period and $\beta = \frac{1}{1+\theta}$ is the rate of time discount, then it can be shown that

$$E_t U_t - U_t^* \simeq -\frac{1}{2}[E_t(y_t - y_t^*)^2 + \lambda(E_t \pi_{t+1} - \pi_t^*)^2] \tag{29}$$

where U_t^* is utility at c_t^* and $y_t^*(= c_t^*)$, the optimal levels of consumption and total output, π_t^* is the target level of inflation and $\lambda = \gamma\left(\frac{\sigma\rho(1-\rho)}{\theta+\rho}\right)^2 \simeq \gamma[\sigma(1-\rho)]^2$ as $\beta \simeq 1$. If expenditure shares

[2] For a critical review of the optimal monetary policy literature, see Svensson (2003).

are approximately constant near equilibrium so that $\sigma = 1$, and the average duration of price changes is 4 months - see Bils, Klenow and Krystov (2003) - so that $\rho = 0.25$ per month, then for quarterly data $\rho \simeq 0.6$. A typical value of γ in calibration exercises is 2. On these grounds we assume that $\lambda \simeq 0.16\gamma \simeq 0.32$. In Table 1 we list the alternative weighting schemes we examine initially.

The alternative specifications of the instantaneous welfare that we consider are summarised in the following table.

Table 1: Alternative specifications of the objective function

	Policy Weights		
Model	λ_π	λ_y	$\lambda_{\Delta rs}$
M1 - flexible inflation targeting	1	1	0
M2 - strict inflation targeting	1	0	0
M3 - strict output targeting	0	1	0
M4 - flexible IT with int smoothing	1	1	1
M5 - flexible IT with DSGE prefs	1	0.32	0

Each specification of the welfare function may be computed for each optimal policy option. In addition, for comparison purposes, we compute monetary policy based on the Taylor rule.

5 Empirical results

5.1 VAR

The VAR model of the economy must contain the target variables and the policy instruments. It may also include other non-policy variables. The choice of non-policy variables determines which shocks we would like to include in our analysis. The vector of variables modelled in the VAR is $\mathbf{z}'_t = (\pi_t, y_t, \pi^o_t, r^{10}_t, r^3_t, R_t)$, where π_t, the rate of inflation, and y_t, the output gap, which is measured by the deviation of log GDP from a quadratic time trend, are the target variables. R_t is the effective federal funds rate and is the monetary policy instrument. Additional non-policy variables are π_t^o, the rate of change of the price of oil, and r_t^{10}, the 10-year long-term interest rate, and r_t^3, the 3-month Treasury bill rate. r_t^{10} and r_t^3 are included to help capture the monetary transmission mechanism via the term structure. The data are quarterly for the US 1964q1-2007q3 and the VAR has a lag length of 8.[3] As we wish to conduct a counter-factual analysis compared with the actual data, we use the sample averages as the target levels.

To check the structural stability of the VAR we computed recursive Chow tests, but with one marginal exception none were significant. We also examined recursive estimates of the VAR coefficients. These showed little variation beyond the initial start-up observations. These findings support Rudebusch's conclusion that a monetary policy VAR for the US does not display much evidence of the sort of structural instability predicted by the Lucas Critique, Lucas (1976).

5.2 Option 1

First we consider the results for option 1 in some detail. We give the outcomes for each welfare function. We focus on the results for optimal policy and do not report the VAR estimates as these are not of especial interest.

5.2.1 Interest rates

In Table 2 we report the policy rules implied by the different specifications of the welfare function for the standard and the PVAR approaches together with federal funds equation in the VAR. We note that, although under the standard approach, executing the optimal interest rate policy makes no difference to the one-period ahead forecasts of inflation and output, it is still of interest to know what the optimal interest rate would be.

[3] The data are taken from the Federal Reserve Bank of St. Louis' database.

Table 2: Interest rate rules, long run coefficients

$$R = \theta_0 + \theta_1\pi + \theta_2 y + \theta_3\pi^o + \theta_4 r^{10} + \theta_5 r^3$$

	VAR	**Standard**					**PVAR**				
		M1	**M2**	**M3**	**M4**	**M5**	**M1**	**M2**	**M3**	**M4**	**M5**
π^c	0.16	0.65	0.86	0.11	0.55	0.76	0.49	0.65	−0.03	0.42	0.57
y	0.09	0.44	0.17	1.20	0.32	0.29	0.04	0.04	−0.03	0.00	0.04
π^o	0.00	−0.06	−0.08	0.00	−0.05	−0.07	−0.06	−0.11	0.05	−0.04	−0.08
r^{10}	0.02	0.11	0.15	−0.06	0.1	0.13	0.19	0.28	−0.12	0.18	0.24
r^3	1.03	0.81	0.72	1.10	0.83	0.76	0.62	0.47	1.04	0.67	0.55
const	−0.37	−1.89	−2.50	−0.01	−1.47	−2.22	−0.78	−1.20	1.28	−0.69	−1.02

We recall that the interest rates r^{10} and r^3 are present in the rules. Since, as a first approximation, they will move one-for-one with the federal funds rate in the long run, in order to obtain the long-run response of the federal funds rate to the other variables, and in particular inflation and the output gap, it is necessary solve the equations on the assumption that the three interest rates are identical. As expected, we find that interest rates respond the most strongly to inflation under strict inflation targeting, M2. Under flexible inflation targeting (M1, M4 and M5) we find that the interest rate responds more strongly to the output gap in the standard than the PVAR method. In fact, under the PVAR method the response under flexible targeting is similar to that under strict inflation targeting.

In Table 3 we report the variability of the federal funds rate for the different welfare functions.

Table 3: Interest rate volatility: standard deviation

	Standard	PVAR
M1	4.01	2.96
M2	4.72	3.52
M3	4.01	2.75
M4	3.44	2.71
M5	4.27	3.19

The standard deviation for the Taylor rule is 4.79 and the actual standard deviation of the Federal Funds rate is 3.27. Thus the standard deviation for the PVAR is less than that for the standard approach. Apart from strict inflation targeting it is also less than the actual standard deviation. In Figures 1-5 we plot the time series behaviour of interest rates under control and compare these with the actual behaviour. In Figure 6 we compare the alternative interest rates under the PVAR approach with the actual values and those based on a Taylor rule.

Figure 1

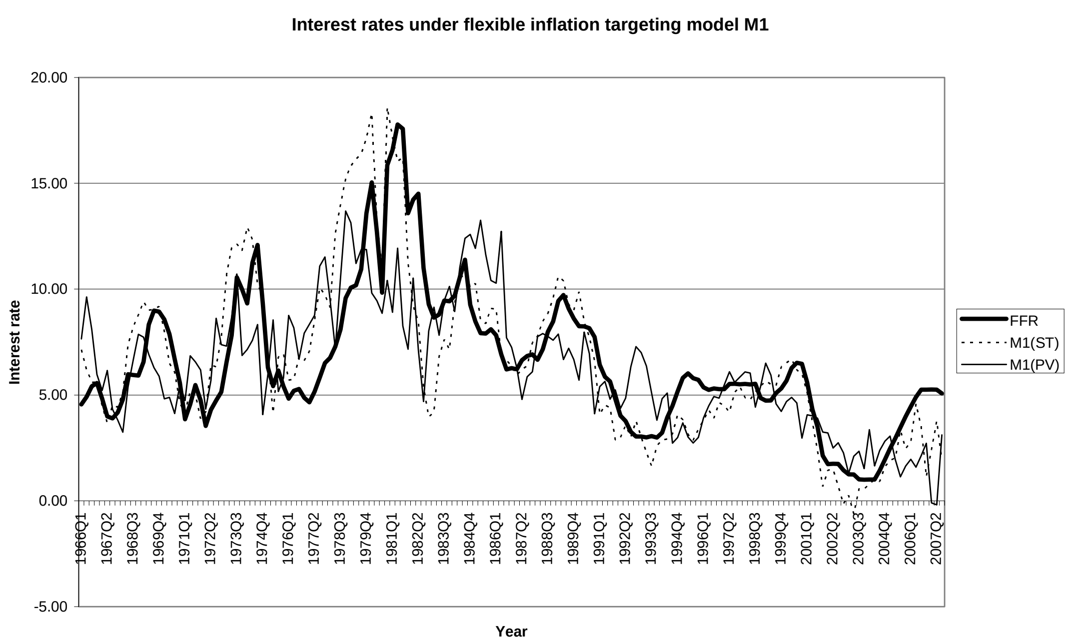

Figure 2

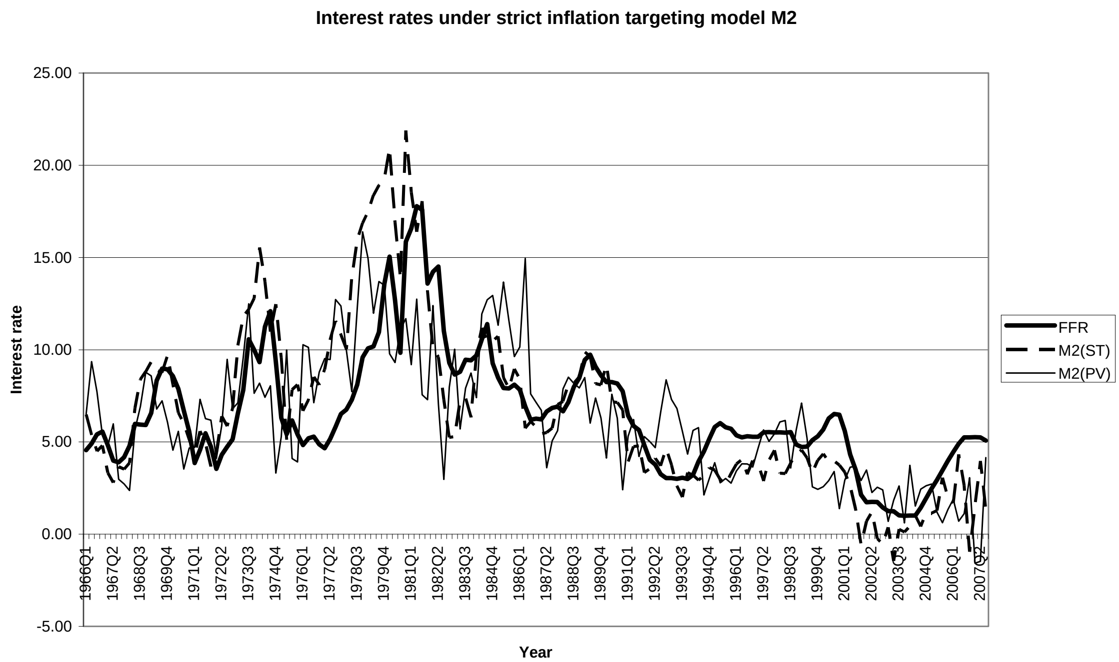

Figure 3

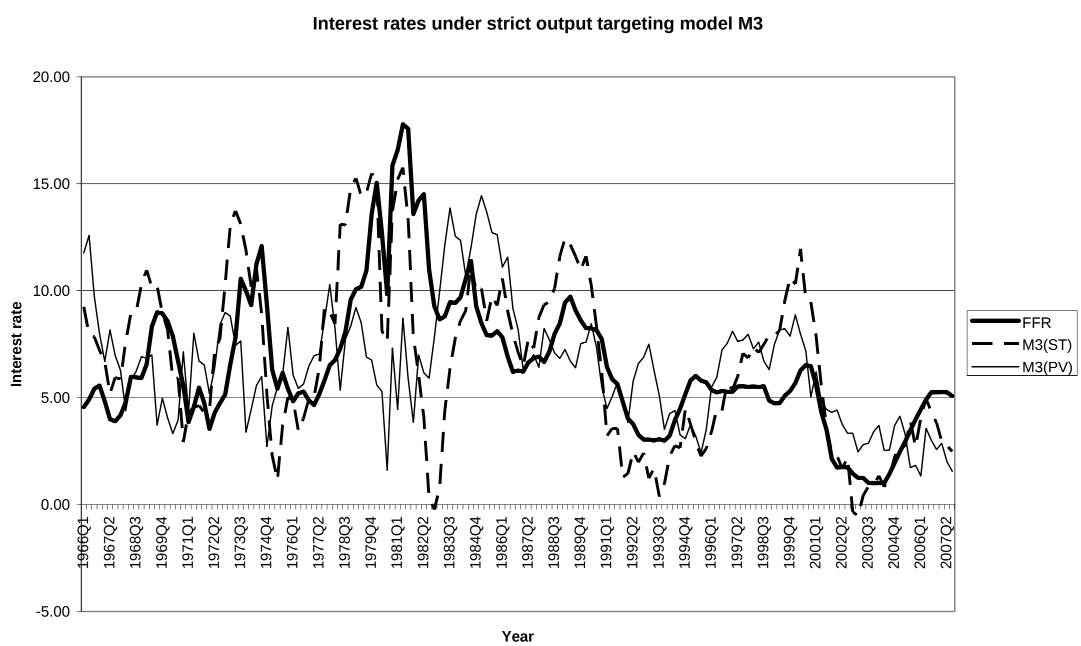

Figure 4

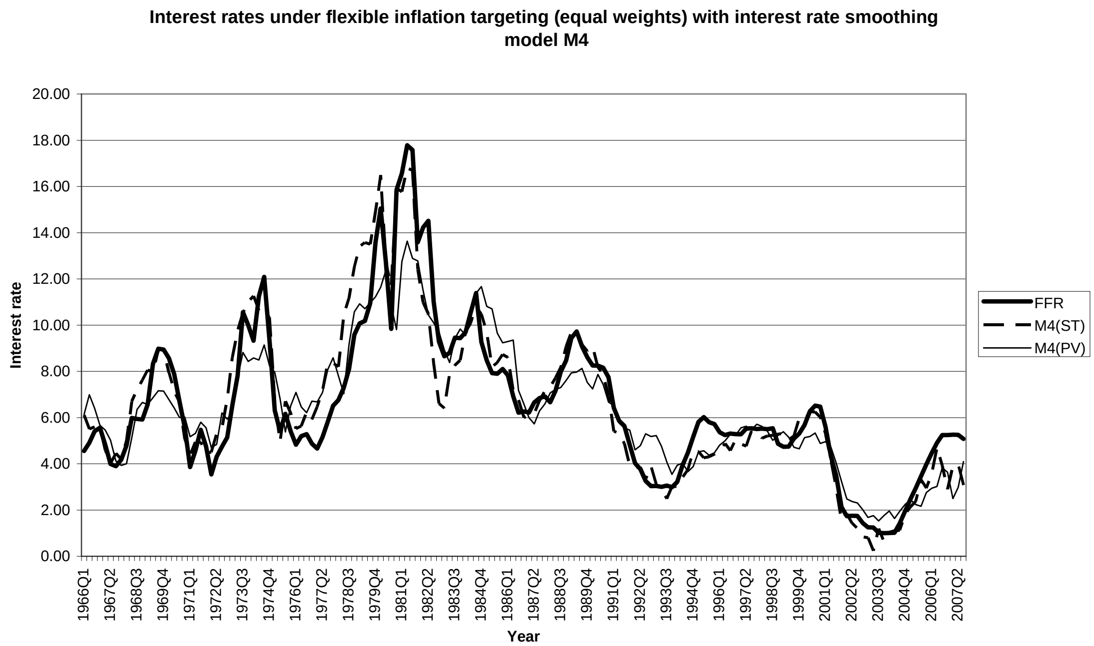

Figure 5

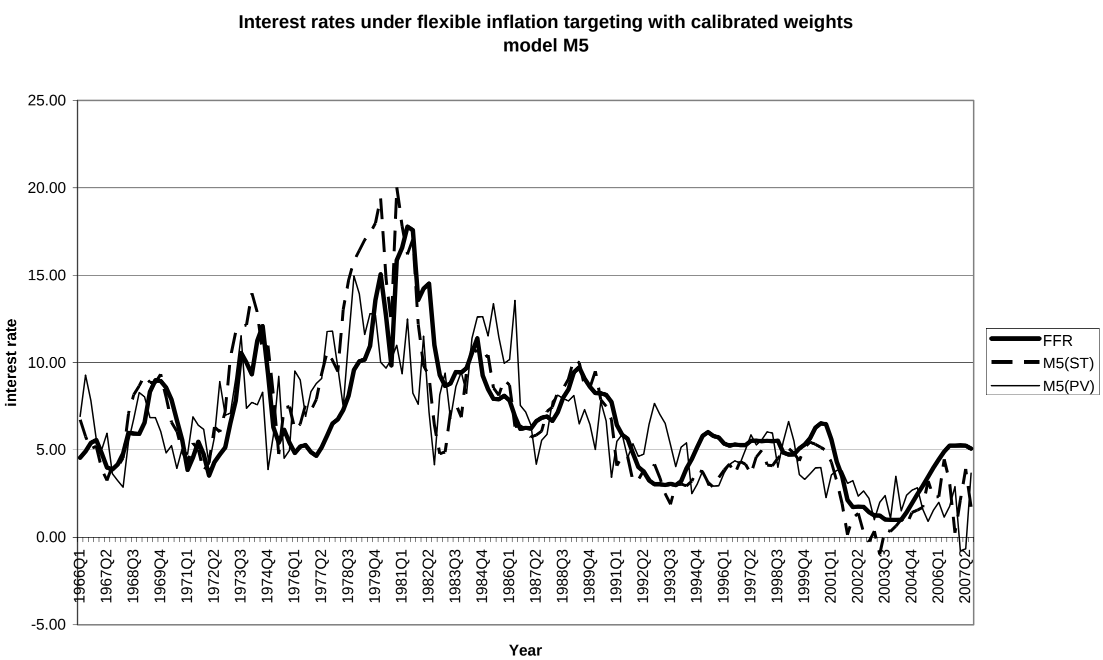

Figure 6

Interest rates for the PVAR under different models and the Tayor rule

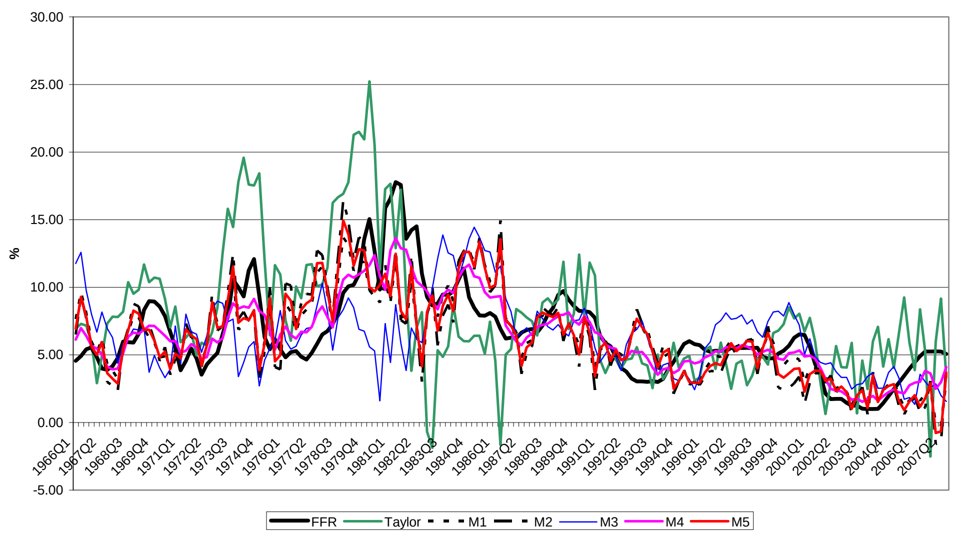

With the exception of the strict output targeting and the Taylor rule, in general, the Federal Funds rate under optimal policy is not hugely different from its actual historic value. Nonetheless, there are episodes when they do differ, especially when the PVAR method is used. For example, in the mid 1980's and early 1990's the optimal PVAR rate implies tighter monetary policy. Of more immediate interest, according to the PVAR method rates should have been higher in 2003-4 but then lower in 2005-6.

In contrast, interest rates based on the Taylor rule are much more volatile than the actual and the PVAR rates in which the control of inflation is an explicit objective. Strict output targeting (M3) results in much lower interest rates over the period 1979-82 and higher rates for the period 1997-2000.

5.2.2 Inflation and output

In order to calculate the outcomes for inflation and output we have explained that we should use the VAR under control. We have also argued that for the standard approach for Option 1 this is simply the one-period ahead forecasts from the original VAR. In order to make a comparison of the standard with the PVAR approach meaningful we start by reporting the inflation and output outcomes based on the state-space representations under control. Since there is no point in reporting the outcomes for the standard approach calculated from the VAR under control, we then just consider the outcomes for the PVAR.

In Table 4 we report the standards deviations for inflation and the output gap for the standard and the PVAR approaches based on the state-space representations under control. The actual standard deviations for inflation and output are 2.97 and 2.13, respectively. The volatility of inflation is lower for the PVAR than the standard approach for each specification of the welfare function but higher for the output gap. It is also lower than the actual volatility of inflation but not the actual volatility of the output gap.

Table 4

	ST		PVAR	
	$sd(\pi_t)$	$sd(y_t)$	$sd(\pi_t)$	$sd(y_t)$
M1	2.79	1.81	2.56	2.36
M2	2.83	1.91	2.77	2.44
M3	2.76	1.58	2.09	2.45
M4	2.76	1.93	2.54	2.10
M5	2.81	1.86	2.66	2.38

In Figure 7 we show the time series behaviour of optimal inflation and the output gap for model M5 for the standard and the PVAR approaches based on the state-space representations under control together with their actual values. The differences between the three series are not large for both inflation and output, nonetheless, the PVAR differs more from the actual values than the standard approach. Similar results occur for the other specifications of the welfare function.

Figure 7

Inflation and output for model M5: actual, standard and PVAR methods

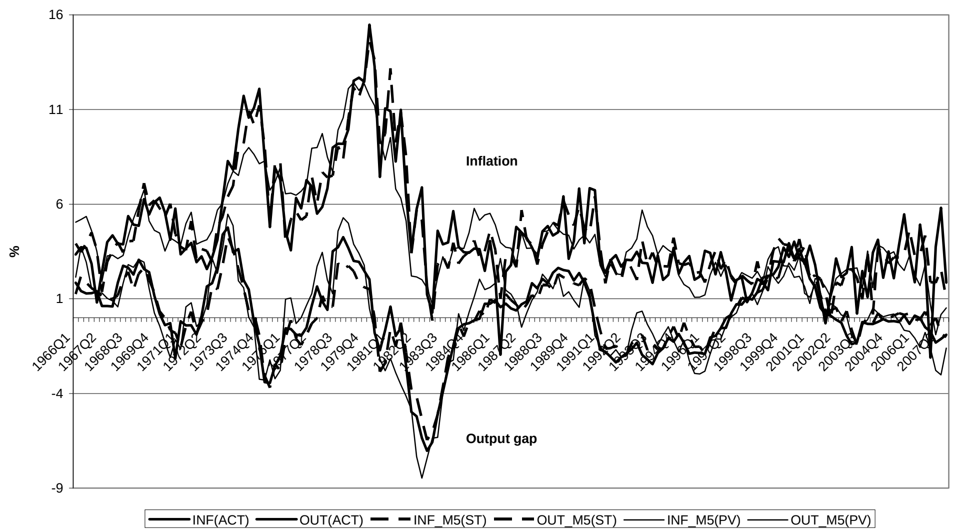

In Figure 8 we show the outcomes for inflation for all five specifications of the welfare function for just the PVAR. Figure 9 gives the corresponding outcomes for the output gap. With the exception of strict output targeting, the outcomes for inflation are fairly similar and they are even closer for the output gap.

Figure 8

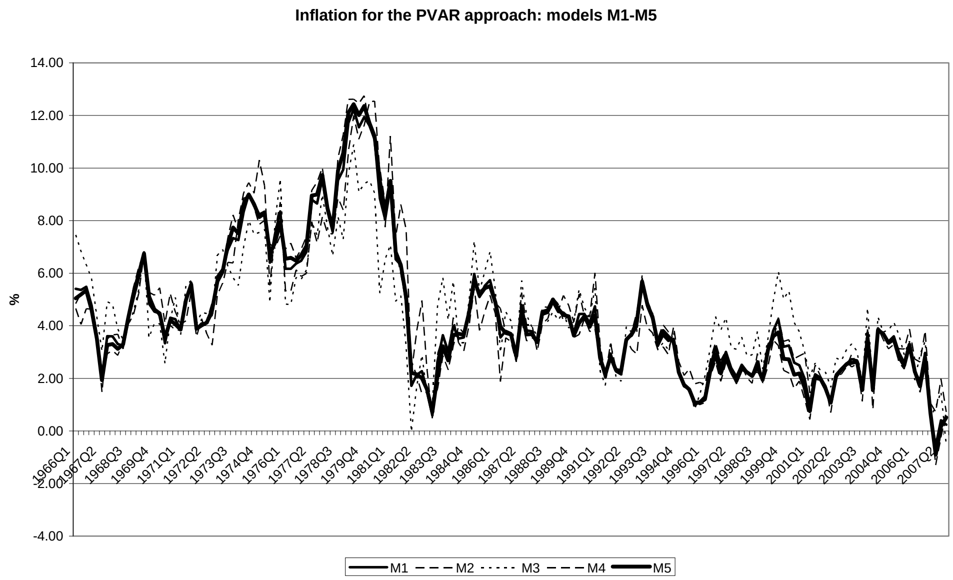

Figure 9

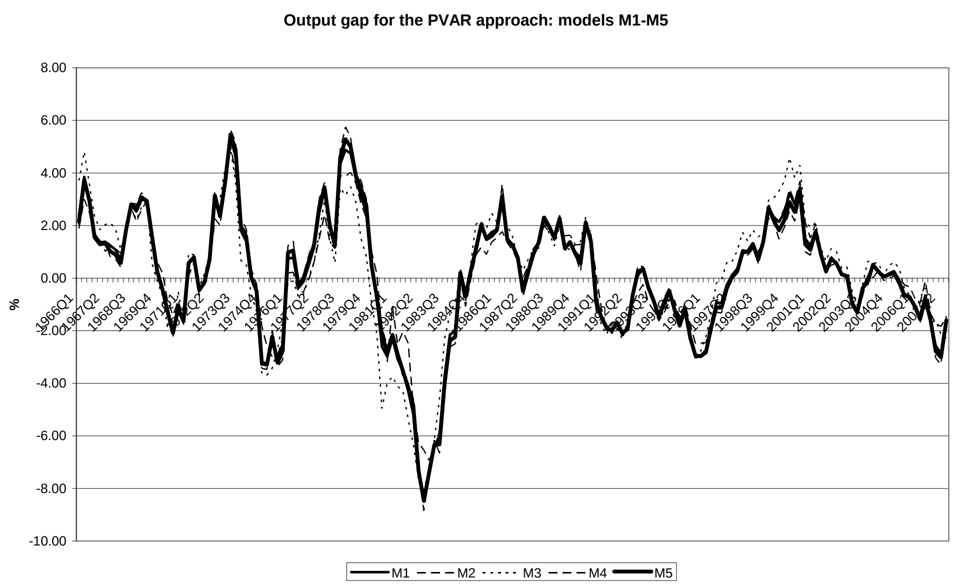

5.2.3 Welfare

In Figure 10 we examine the inflation-output trade-off in more detail by considering the effect of different choices of the welfare function through using a continuum of relative weights on inflation and the output gap. We report the results with and without interest rate smoothing. We use the state-space representation of the VAR. Compared with the standard method, the variability of inflation is generally lower using the PVAR method but that of the output gap is higher. Smoothing interest rates seems to improve the trade-off for the PVAR method principally by lowering the variability of the output gap

Figure 10. Inflation-output trade off

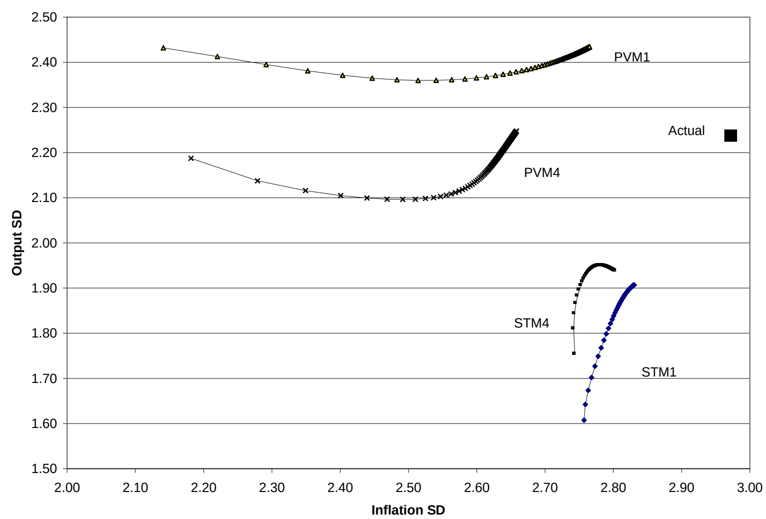

5.3 Options 1-4

We now examine the outcomes for the four options for just the PVAR method for the wefare function M5 over the period 1984-2007. We do not consider the standard approach as without building back the disturbances it gives the same results as the one-period forecasts from the original VAR, and when building back the disturbances it gives the orginal data. Option 1, previously analysed, involves forecasting one period ahead without building back the errors we know to have occurred. In option 2 we include the errors. In options 3 and 4 we re-estimate the VAR each period before calculating the optimal policies corresponding to options 1 and 2; we also report the actual historic values. We recall that for the PVAR the optimal Federal Funds rate will be the same under options 1 and 2 and under options 3 and 4, but the outcomes for inflation and output will differ.

5.3.1 Interest rates

The optimal Federal Finds rates are shown in Figure 11. They show considerable differences. The optimal rates are more volatile than the actual rates, and re-estimating the VAR each period results in greater volatility than using the VAR estimates for the whole sample. This is consistent with needing to have a more flexible monetary policy the more one wishes to stabilise inflation and output as the interest rate must absorb more of the shocks to inflation and output. And since re-estimating the VAR each period will mean the VAR coefficients will vary instead of staying fixed, even with the same shocks to inflation and output, interest rates will fluctuate more.

Figure 11

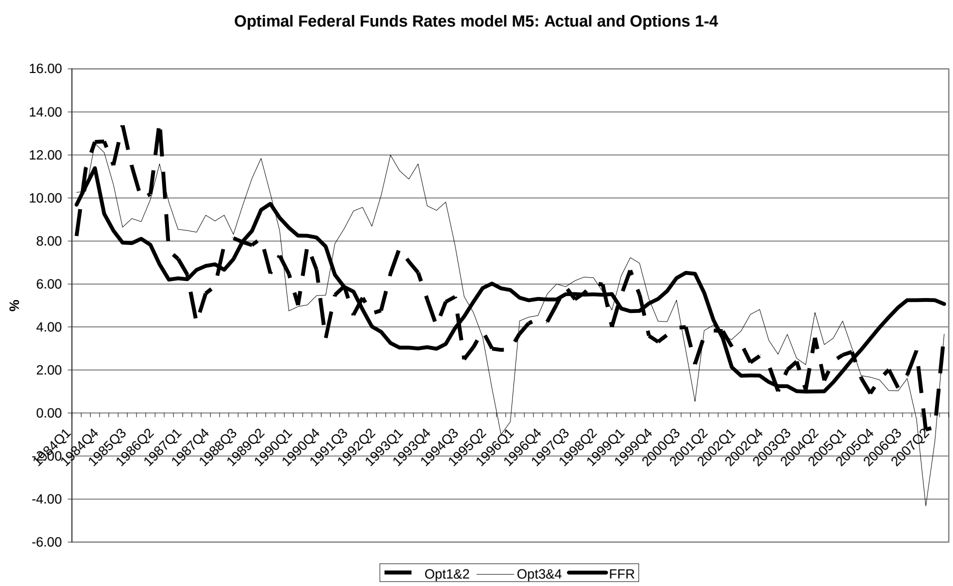

5.3.2 Inflation and output

We now consider the consequences for inflation and output over the period 1984-2007. In Table 6 we report the standard deviations for inflation and output for the PVAR approach for the four options and the five welfare specifications. Even though the optimal interest rates are the same for options 1 and 2 and for options 3 and 4, the inflation and output gaps will be different. The table shows that re-estimating the VAR each period tends to increase the volatility of inflation but not output. Adding back the disturbances has little effect.

Table 6. Standard deviations for inflation and the output gap 1984.1-2007.3

Option	$M1$		$M2$		$M3$		$M4$		$M5$	
	π	y	π	y	π	y	π	y	π	y
1	2.66	0.16	1.76	0.75	4.41	0.98	2.54	0.24	2.22	0.45
2	3.10	0.14	2.19	0.73	4.84	1.00	2.97	0.22	2.65	0.43
3	3.71	0.45	2.80	0.22	4.81	1.25	3.24	0.12	3.43	0.25
4	3.62	0.25	2.72	0.42	4.73	1.05	3.16	0.08	3.35	0.05

Standard deviations for the observed data: Inflation: 2.67; Output: 0.41

In Figures 12 and 13 we show the time series behaviours of inflation and the output gap for the welfare specification M5. From 1996 inflation is similar for the different options and similar to actual inflation, but prior to this optimal inflation tends to be higher than actual inflation but to differ between the options. The differences are most marked in the periods 1991-4, 2001-4 and from mid 2006. The explanation for the differences in the period 1991-4 is to be found in the output gap series as during this period the actual output gap was negative whilst as a result of stimulating inflation it becomes positive for options 3 and 4 and closer to zero under options 1 and 2. In the period 2001-4 output would have been higher without inflation being much affected. And from mid 2006 the optimal output gap is much more negative than actually occurred. This is because optimal policy was trying to bring down the high actual inflation over this period. Another part from these periods the output gap is similar for all of the options and similar to the actual output gap.

Figure 12

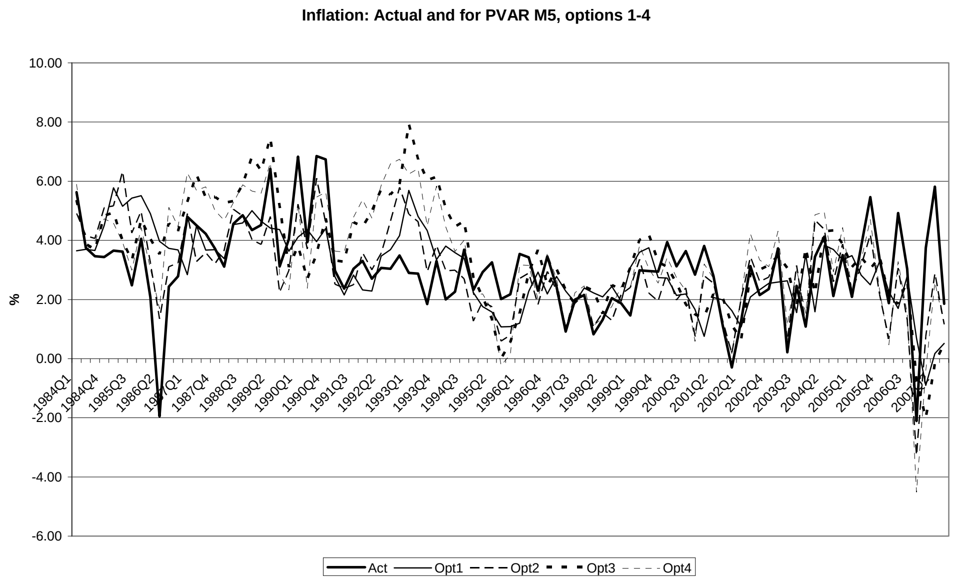

Figure 13

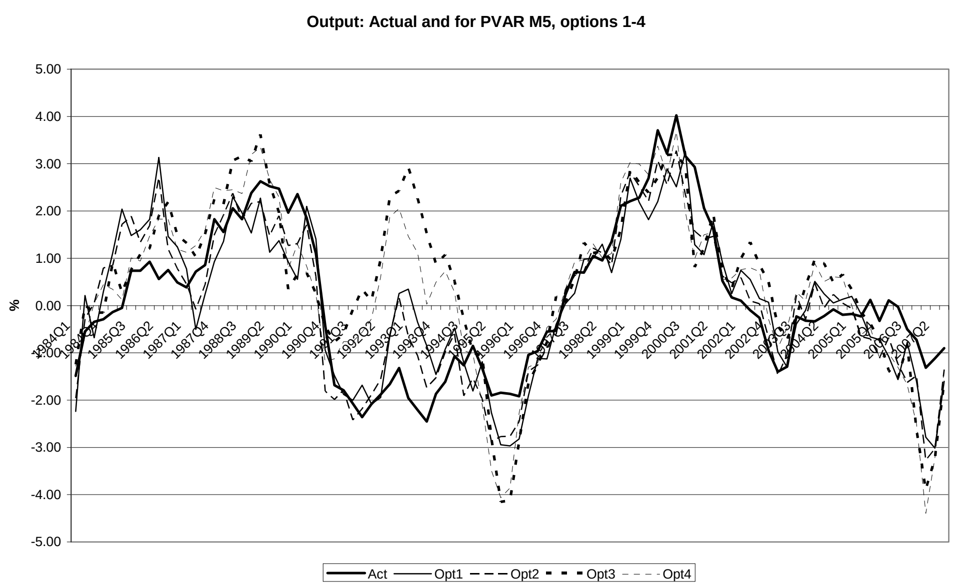

6 Conclusions

In this paper we have suggested a way of formulating optimal policy based on a VAR that avoids many of the problems found in the standard approach. For example, and perhaps the most important advantage, our proposed PVAR method does not involve having to make any identifying restrictions in the VAR. It is suitable for re-optimising policy each period based on up-dated estimates of the VAR as one would do in real-time applications. The optimal solutions for the state vector are the same whether one uses a state-space representation of the VAR under control or a VAR based on the original variables. None of these holds for the standard approach. Since the whole process is easily automated, the PVAR method may provide a useful benchmark for use in real time against which to compare other, probably far more labour intensive, policy choices.

Although basing optimal policy on a VAR has the merit of simplicity, it is not without its drawbacks. We have shown that as a result of implementing optimal policy, the VAR under control is different from the original VAR. This is not necessarily a problem in itself, but it does draw attention to the fact that any previous changes of policy are likely to have caused structural change in the original VAR. This shows the vulnerability - at least in theory - of any VAR to structural change. The problem is further exacerbated because the VAR is just a particular time series representation of a structural model. If the parameters of the structural model alter as a result of policy changes, then we would expect the VAR coefficients to change too. In practice, like Rudebusch, we find little evidence of structural change in the dynamics of a VAR suitable for analysing monetary policy fo the US.

Another drawback of using a VAR is that it is not suitable for handling the effects on non-policy variables of anticipated policy changes. One cannot avoid using a structural rational expectations model if one wishes to analyse this problem. To avoid any misapprehensions, therefore, we emphasise that in arguing the merits of adopting the PVAR method for formulating policy based on a VAR, we are not suggesting that using a VAR is necessarily preferable to using a well specified structural model.

We illustrate the use of the PVAR method by analysing monetary policy for the US since 1966. We examine the effect of different specifications of the welfare function and different ways of implementing optimal policy. And we compare the PVAR method with the standard approach and with a Taylor rule. Our results suggest that optimal monetary policy obtained using the PVAR method would have been tighter during the periods 1991-4 and 2001-4 and looser from 2005. As a result, output would have been higher over the periods 1991-4 and 2001-4, but lower from 2005. Inflation is little different, except from 2005 when it would have been lower.

References

Barro, R. J. and D.B. Gordon (1983). "Rules, discretion and reputation in a model of monetary policy", *Journal of Monetary Economics*, 12, 101-121.

Bellman, R. E. (1957). *Dynamic Programming*, Princeton University Press, Princeton.

Bernanke, B. S. and A.S. Blinder (1992). "The Federal Funds Rate and the Channels of Monetary Transmission", *The American Economic Review*, 82, 4, 901-921.

Bernanke, B. S. and I. Mihov (1998). "Measuring Monetary Policy", Quarterly Journal of Economics, 113:3, 869-902.

Bils, M., P.J. Klenow and O. Kryvtsov (2003), Sticky Prices and Monetary Policy Shocks," *Federal Reserve Bank of Minneapolis Quarterly Review*, 27, 2-9.

Clarida, R., J. Gali and M. Gertler (1999). "The Science of Monetary Policy: A New Keynesian Perspective", *Journal of Economic Literature*, 37:4, 1661-707.

Chow, G. C. (1976). *Analysis and Control of Dynamic Economic Systems*, Princeton.

Christiano, L. J., M. Eichenbaum and C.L. Evans (1999). "Monetary Policy Shocks: What Have We Learned and to What End?," in J. B. Taylor and M. Woodford, eds., *Handbook of Macroeconomics*, Vol. 1A, 65-148, Amsterdam: Elsevier.

Goodhart, C. A. E. (1999). "Central Bankers and Uncertainty", *Proceedings of the British Academy*, 101, 229-271.

Kydland, F. E. and E. Prescott (1977). "Rules Rather than Discretion: The Inconsistency of the Optimal Plan", *Journal of Political Economy*, 85, 3, 473-492.

Leeper, E.M. C.A. Sims, T. Zha, R.E. Hall and B.S. Bernanke (1996), "What Does Monetary Policy Do?" *Brookings Papers in Economic Activity*, 1-78.

Ljungqvist, L. and T.J. Sargent (2004). *Recursive Macroeconomic Theory*, MIT.

Lucas, R. (1976). "Econometric Policy Evaluation: A Critique", *Carnegie-Rochester Conference Series on Public Policy*, 1, 19–46.

Martin, B, and C. Salmon (1999). "Should Uncertain Monetary Policy-Makers Do Less?", Bank of England Working Paper No. 99.

Rogoff, K. (1985). "The Optimal Degree of Commitment to an Intermediate Monetary Target", *Quarterly Journal of Economics*, 100, 4, 1169-1189.

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

Rudebusch, G.D., (1998). "Do Measures of Monetary Policy in a Var Make Sense? *International Economic Review*, 39, 907-931.

Rudebush, G. D. (2001). "Is The Fed Too Timid? Monetary Policy in an Uncertain World", *The Review of Economics and Statistics*, 83, 203–17.

Rudebusch, G.D. (2002) "Assessing the Lucas Critique in monetary policy models", FRB of San Francisco Working Paper 2002-02.

Rudebusch, G. D., and L.E.O. Svensson (1999). "Policy Rules for Inflation Targeting," in Monetary Policy Rules. John B. Taylor, ed. Chicago: Chicago U. Press, 203—246.

Sack, B. (2000). "Does the Fed Act Gradually? A VAR Analysis", *Journal of Monetary Economics*, ???.

Sack, B., and Wieland, V. (2000). "Interest Rate Smoothing and Optimal Monetary Policy: A Review of Recent Empirical Evidence", *Journal of Economics and Business*, 52, 205-228.

Sims, C.A. (1980), "Macroeconomics and Reality," *Econometrica,* 48, 1-48.

Stock,J. and M.W. Watson (2001), "Vector autoregressions", Journal of Economic Perspectives, 39, 101-116.

Svensson, L.E.O., (2003). "What Is Wrong with Taylor Rules? Using Judgment in Monetary Policy through Targeting Rules", *Journal of Economic Literature*, 41, 2, 426-477.

Svensson, L.E.O., (1997). "Optimal Inflation Targets, "Conservative" Central Banks, and Linear Inflation Contracts", *American Economic Review*, 87, 1, 98-114.

Taylor, J. B. (1993). "Discretion versus Policy Rules in Practice." Carnegie-Rochester Conference Series on Public Policy 39, 195-214.

Taylor, J. B. (1995). "Monetary Policy Implications of Greater Fiscal Discipline", Symposium Proceedings, *Budget Deficit and Debt: Issues and Options*, Federal Reserve Bank of Kansas City.

Taylor, J.B. (2000). "The Policy rule Mix: A Macroeconomic Policy Evaluation", in Calvo, G, Dornbusch, R., Obstfeld, M., eds., *Money, Capital Mobility and Trade, Essays in Honour of Robert Mundell*, MIT Press, 505-518.

Walsh, C. E. (2003). *Monetary Theory and Policy,* 2nd ed, MIT Press.

Wickens,M.R (1996) "Interpreting cointegrating vectors and common stochastic trends", Journal of Econometrics, 74, 255-271.

Wickens,M.R. (2008). *Macroeconomic Theory: A Dynamic General Equilibrium Approach*, Princeton: Princeton University Press.

Wickens, M.R. and R. Motto (2001) "Estimating Shocks and Impulse Response Functions", *Journal of Applied Econometrics*, 16, 371-387.

Woodford, M. (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton: Princeton University Press.

Woodford, M. (2007). "The case for forecast targeting as a monetary policy strategy", *Journal of Economic Perspectives*, 21, 3-24.