BANK OF ENGLAND

CENTRE FOR CENTRAL BANKING STUDIES

ECONOMIC MODELLING AND FORECASTING

Estimating the implicit inflation target of the South African Reserve Bank using state-space models and the Kalman filter

by

Ole Rummel
Centre for Central Banking Studies
Bank of England

9 February 2015

ole.rummel@bankofengland.co.uk   1

1   Introduction

Many countries do not announce an (implicit) inflation objective. Prior to 25 January 2012, for example, the US was the most prominent example of a country not having either an implicit or explicit quantitative inflation objective. On that day, the Federal Reserve, in its first ever 'longer-run goals and monetary policy strategy' statement, said that an inflation rate of 2 per cent was most consistent over the longer run with its congressionally mandated goals of price stability and full employment.

South Africa, the subject of this exercise, adopted inflation targeting (IT) as a formal framework for monetary policy in February 2000. Starting in 2002, an inflation target was specified on the basis of an annual inflation rate of between 3 and 6 per cent. This range has been largely maintained since then, although the target measure was re-defined several times. At the same time, the South African Reserve Bank (SARB) has refrained from announcing a mid-point. In this regard, South Africa stands out amongst its international peers in using an inflation target range that exceeds a 1 percentage point band and not having an official point inflation target. This set-up has advantages as well as disadvantages. For the monetary policymaker, the lack of transparency regarding an inflation target mid-point provides greater flexibility in implementing monetary policy in an environment of volatile inflation. But this flexibility comes at a cost, which is that it may undermine the central bank's efforts to anchor medium- and long-term inflation expectations and may complicate market participants' projections regarding the future path of the policy rate and the monetary policy stance more generally.

The absence of official quantitative inflation objectives has not stopped economists from trying to estimate implicit or unobserved inflation objectives empirically. Many estimates of the implicit inflation target rely on the assumption that it remains constant over the period under observation and analysis. A prominent example is Clarida *et al.* (1998), who estimated the Federal Reserve's policy reaction function over the period from 1979 to 1994 using a model with a constant inflation target. Their (fixed) estimate of this target for the US was equal to 4 per cent.

But there are a number of reasons to believe that a central bank's (implicit) inflation target may vary over time. A time-varying parameter model and use of the Kalman filter (Kalman (1960, 1963) and Kalman and Bucy (1961)) allow us to relax this assumption and estimate a time-varying implicit inflation target. One of the advantages of the Kalman filter is that no assumptions about whether or how the implicit target has varied are necessary.

The aim of this exercise is to replicate the results in Klein (2012), who was interested in 'revealing' the SARB's *ex post* point inflation target and assess to what extent it has varied since the introduction of IT.[1] The underlying methodology itself is based on Leigh (2008).[2]

2   Specification and estimation of monetary policy reaction functions

We first need to give some thought on how to ascertain an implicit inflation objective. There exists a large literature on monetary policy rules and on how monetary policy-makers have reacted to deviations of inflation and real activity from their – explicit or implicit – target levels. This has become enshrined in the concept of a Taylor rule.

---

[1] Many thanks to Nir Klein for furnishing me with the original data used in his study.
[2] Leigh (2008) estimated the Federal Reserve's implicit inflation target using a time-varying parameter (state-space) model and the Kalman filter.

While there are several ways of describing monetary policy, most new Keynesian (NK) macro models nowadays are closed using a Taylor rule, which is meant to represent the behaviour and preferences of the monetary policy maker. It is worth pointing out that while no central bank has formally declared that it is following an interest rate rule such as the Taylor rule, the latter may still provide a parsimonious description of actual policy. A typical (first-generation) Taylor rule looks as follows:

$$i_t^* = r^n + \pi_t^e + (\beta - 1)(\pi_t^e - \pi^*) + \gamma \tilde{y}_t \tag{1}$$

where i_t^* is the target level of the short-term monetary policy instrument in time period t (here, the SARB's repo rate), π_t^e is expected inflation (which could be the inflation forecast), $\tilde{y}_t$ is the output gap (the percentage difference between actual and potential real GDP), r^n is the natural rate of interest (defined as the equilibrium real rate of return in the case of fully flexible prices) and π^* is the inflation target or a desired inflation rate.[3] The Taylor rule describes the behaviour and preferences of the monetary authority. In the Taylor rule, the policy instrument – generally, but not always, a (short-term) interest rate – responds to economic variables with the aim of achieving a pre-specified policy objective, such as hitting an inflation, exchange-rate or money-supply target. In the case of the Taylor rule given by equation (1), the central bank adjusts the nominal interest rate in response to deviations from the target values for inflation and output.[4] All the parameters in the model are assumed to be positive.

But most empirical studies of Taylor rule variants have found that the lagged nominal interest rate enters equation (1) and that it is strongly significant. In practical terms, this equates to central bankers adjusting interest rates more cautiously than predicted by standard models. Interest-rate smoothing was therefore introduced to make the dynamics of i_t more realistic.[5] The (generalised) empirical Taylor rule in equation (1) thus becomes:

$$i_t^* = r_t^n + \pi_t^e + (\beta - 1)(\pi_t^e - \pi_t^*) + \gamma \tilde{y}_t \tag{2}$$

$$i_t = (1 - \rho)i_t^* + \rho i_{t-1} + \varepsilon_t \qquad 0 \le \rho < 1 \tag{3}$$

or, combining equations (2) and (3):

$$i_t = (1 - \rho)(r_t^n + \pi_t^e + (\beta - 1)(\pi_t^e - \pi_t^*) + \gamma \tilde{y}_t) + \rho i_{t-1} + \varepsilon_t \tag{4}$$

with $\beta > 1$ and $0 \le \rho < 1$ denoting the degree of lagged dependence in the interest rate (i.e., interest-rate smoothing). If ρ is high, the central bank will normally adjust the policy stance only gradually in response to economic developments or prospects.

The restriction that $\beta > 1$ should hold has become enshrined as the 'Taylor principle'. Under such parameter values, the Taylor rule will generally ensure a unique and deterministic rational expectations equilibrium, as shown by Woodford (2003, pages 252-61). In other

---

[3] In this exercise we shall not consider putting the exchange rate in the Taylor rule, as is done in Lubik and Schorfheide (2007).
[4] An influential empirical paper in this respect is Clarida *et al.* (1998), while the theory can be found in Taylor (1999).
[5] Note, however, that Woodford (2003, p. 280) has also provided a theoretical rationale for the inclusion of the lagged interest rate in equation (4).

words, this condition stabilises inflation: when the inflation forecast rises above target, the policymaker raises nominal interest rates enough to raise the target for the real policy rate. Variations in the Taylor rule are used to explain monetary policy across different regimes.

For a number of reasons, the Taylor rule has become a cornerstone of modern monetary policy analysis. To begin with, it is simple to understand and transparent. Second, it reflects the move to 'models without money', in which the interest rate – rather than a (narrow) monetary aggregate – has become the main policy instrument.

2.1   The data

The EViews file **implicit_target.wf1** contains data for South Africa, including the SARB's repo rate as the dependent variable. The analysis uses quarterly data over the period from 2001 Q1 to 2011 Q4. Most of the data are available from the SARB's website (www.resbank.co.za), including the Reuters inflation expectations survey data discussed below. To open the EViews workfile from within EViews, choose **File**, **Open**, **EViews Workfile…**, select **implicit_target.wf1** from its current location, update the default directory by checking the small box and click **Open**. Alternatively, you can double-click on the workfile icon outside of EViews, which will open EViews automatically.

There are a number of inherently unobservable time series that enter the model as independent variables on the right-hand side, and Klein either has to make (reasonable) assumptions about their values or derive some proxy measures indirectly. In particular, the three unobservable variables that enter into equation (4) are:

- the output gap;
- the natural real rate of interest; and
- inflation expectations

2.1.1 Measuring the output gap

The first unobservable variable is the output gap, $\tilde{y}_t$. The unobservable output gap (denoted `gap100`) over the period from 1980 Q1 to 2011 Q4 was computed from the natural logarithm of real GDP (`lnrgdp`), which has been de-trended using the Hodrick-Prescott (HP) filter. This procedure has already been carried out for you, so there is no need to calculate the Hodrick-Prescott filtered series – as well as the output gap – again. Just in case, these are the steps that are required for detrending the data with the Hodrick-Prescott filter. To smooth the series `lnrgdp` using the HP filter, double-click on it to open the series and choose **Proc/Hodrick-Prescott Filter…**, which opens up the Hodrick-Prescott Filter dialog box. We can provide a name for both the smoothed and the cycle series. EViews will suggest a name for the smoothed series, but I have chosen to call it `hp_lnrgdp`. EViews has recognised the fact that the data are quarterly, so has specified a value of 1600 for the smoothing parameter, λ, automatically. You may specify any other value for λ directly or use the frequency power rule of Ravn and Uhlig (2002). Upon clicking **OK**, EViews will display a graph of the filtered series together with the original series. We also note that the workfile now contains the additional series that we have created.

Alternatively, the EViews command when using the command line is either:

```
lnrgdp.hpf(1600) hp_lnrgdp @ gap
```

when you want to extract both the smoothed and the cycle series (which will be given the name of `hp_lnrgdp` and `gap` in the workfile respectively) or simply:

$$\texttt{lnrgdp.hpf(1600) hp\_lnrgdp}$$

when you only want the smoothed series, after which you obviously need to define the output gap as the difference between the two series:

$$\texttt{gap = lnrgdp - hp\_lnrgdp}$$

`gap100`, finally, is equal to 100 times the gap measure just calculated (`gap100 = 100*gap`). The advantage of the HP-filter is that it implicitly allows potential output growth to vary over time. But it is worth pointing out that the Hodrick-Prescott filter in EViews is generally (erroneously) employed as a one-sided filter, meaning that it is used with historical data up to time t (today). It is, however, the case that the HP-filter uses observations at time $t + i$, $i > 0$ to construct the current time point t. Stopping the sample at the current date means that only current and past states, rather than past, current **and** future states influence the current observation. In order to get around this problem, Klein (2012) followed what is now considered to be standard procedure and augmented his real GDP data for South Africa, which end in 2011 Q4, with forecast values from the IMF *World Economic Outlook* database, which take the total sample of the data out to 2017 Q4. By having data beyond the end of the sample period of actual data, the endpoint problems associated with the HP-filter are avoided (Mise *et al.* (2005)).

Figure 1: South African output gap (calculated using the Hodrick-Prescott filter)
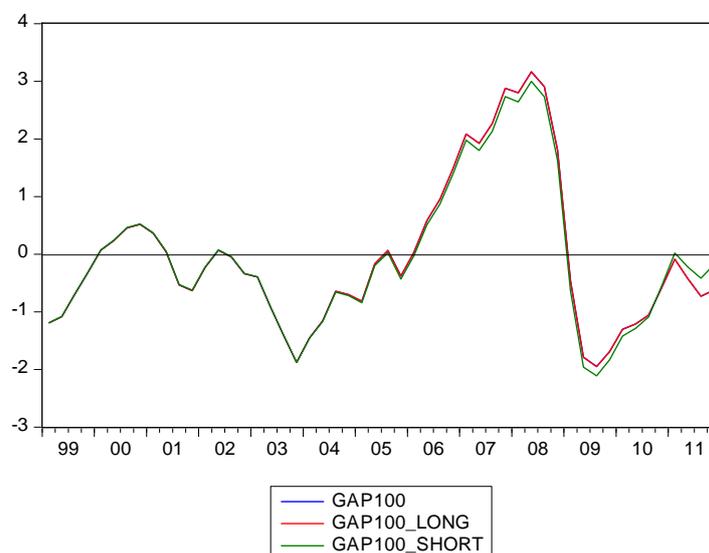

Figure 1 shows that the output gap measure has been quite volatile, fluctuating from -2 per cent in 2009 Q3 to 3.2 per cent in 2008 Q2. We can see that the estimated output gap was still negative at the end of the sample period in 2011 Q4. Moreover, the deviation of the HP-filter derived output gap based on the short sample (1980Q1 to 2011Q4, green line) from the derived output gap using the forecasted values out to 2017Q4 (the red line) becomes quite apparent in the second half of the sample. In particular, the estimated output gap for 2011 Q4

is three times as large using the longer sample period (-0.6 per cent) than with the shorter sample period (-0.2 per cent).

6.2 Measuring the natural real rate of interest

The second unobservable variable in this exercise is the natural real rate of interest, r^n. Klein assumed that the natural real rate of interest was constant over time and equal to its historical average over the sample period, which was about 3.25 per cent.[6]

2.1.3 Measuring inflation expectations

The final unobserved series is inflation expectations, π_t^e. These data are based on a survey that Reuters undertakes on behalf of the SARB and which is available on the SARB's website. In particular, the exercise uses the Reuters Econometer Consensus survey of inflation expectations for four quarters ahead (denoted `exp_reuters1_s` in the workfile). As pointed out in the paper, the advantage of using inflation expectations instead of actual inflation is that it injects a forward-looking element into the Taylor estimation.

2.2 Estimation of monetary policy reaction functions

Having specified the monetary policy rules and either generated or obtained the required data, we can now turn our attention to estimating the Taylor rules we have defined above.

2.2.1 Constant inflation target and natural real rate of interest

As encapsulated in equation (1), the conventional approach to estimating Taylor rules assumes a constant inflation target, $\pi_t^* = \pi^*$, and a constant natural real rate of interest, $r_t^n = r^n$. Equation (1) can thus be rewritten as:

$$i_t^* = \alpha + \beta \pi_t^e + \gamma \tilde{y}_t \qquad (5)$$

where all the variables in equation (5) are observable and the composite intercept term, $\alpha = r^n - (\beta - 1)\pi^*$, consist of both the constant inflation target and the constant natural rate of interest. By estimating α and β and making an assumption regarding the value of r^n, we can solve for the estimated constant inflation target, π^*, using the following expression:

$$\pi^* = \frac{r^n - \alpha}{\beta - 1} \qquad (6)$$

Table 1 presents the estimation results from an OLS regression of equation (5).

---

[6] Klein also points out that this level of the natural real rate of interest is broadly in line with the average estimated potential output growth rate during the period after the end of Apartheid in 1994. The mean of `dhp_lnrgdp_long` (defined as `dhp_lnrgdp_long = 400*d(hp_lnrgdp_long)` is equal to 3.15 per cent over the period from 1994 Q1 to 2011 Q4.

Table 1: Estimation results of Taylor rule with constant inflation target and natural real rate of interest (2001 Q1 – 2011 Q4)

Dependent Variable: RATE
Method: Least Squares
Sample (adjusted): 2001Q1 2011Q4
Included observations: 44 after adjustments

Variable	Coefficient	Std. Error	t-Statistic	Prob.
C	1.688310	3.608356	0.467889	0.6423
EXP_REUTERS1_S	1.306589	0.652047	2.003828	0.0517
GAP100	0.722637	0.237455	3.043261	0.0041

R-squared	0.265687	Mean dependent var	8.886364
Adjusted R-squared	0.229867	S.D. dependent var	2.449166
S.E. of regression	2.149320	Akaike info criterion	4.433926
Sum squared resid	189.4026	Schwarz criterion	4.555575
Log likelihood	-94.54637	Hannan-Quinn criter.	4.479039
F-statistic	7.417269	Durbin-Watson stat	0.144124
Prob(F-statistic)	0.001781		

If we assume that r^n equals the average historical value of the real interest rate, equal to 3.25 per cent, we can solve $\alpha = r^n - (\beta - 1)\pi^*$ for $\pi^* = (r^n - \alpha)/(\beta - 1) = (3.25 - 1.69)/(1.31 - 1) = 5.09$. After estimation, we can get EViews to do the calculation for us. In the command window, type:

```
scalar pi_star = (3.25 - c(1))/(c(2)-1)
```

Double-clicking on `pi_star` in the workfile shows that $\pi^* = 5.09$, which happens to fall within the upper segment of the 3 to 6 percent band announced by the SARB.

Now that we have estimates for $r^n = 3.25$ and $\pi^* = 5.09$, we can go a step further and estimate equation (4) as well. In this, we follow Leigh (2008), who suggested estimating the monetary policy reaction function using ordinary last squares (OLS) before moving to the time-varying parameter model and using maximum-likelihood (ML) estimation. Mirroring the approach in Hamilton (1994), the benefit of doing so is that the parameters to be estimated by ML *via* the Kalman filter in the state-space representation can be initialised using the respective OLS estimates. As such, ML estimation in Section 4.2 below will start from the values of the parameters estimated using OLS and the reaction function with a constant inflation target and a constant natural real rate of interest.

The equation to be estimated is:

$$i_t = (1 - \delta)(r^n + \pi_t^e + (\beta - 1)(\pi_t^e - \pi^*) + \gamma\tilde{y}_t) + \delta i_{t-1} + \varepsilon_t \tag{7}$$

and the estimation results are given in Table 2.

Table 2: Estimation results of Taylor rule with constant inflation target and natural real rate of interest (2001 Q1 – 2011 Q4)

Dependent Variable: RATE
Method: Least Squares
Sample: 2001Q1 2011Q4
Included observations: 44
Convergence achieved after 3 iterations
RATE = (1 - C(1))*(3.25 + EXP_REUTERS1_S + (C(2) - 1)
    *(EXP_REUTERS1_S - 5.09) + C(3)*GAP100) + C(1)*RATE(-1)

	Coefficient	Std. Error	t-Statistic	Prob.
C(1)	0.858711	0.047437	18.10208	0.0000
C(2)	1.592265	1.178960	1.350567	0.1842
C(3)	2.549162	0.906819	2.811102	0.0075
R-squared	0.918340	Mean dependent var		8.886364
Adjusted R-squared	0.914357	S.D. dependent var		2.449166
S.E. of regression	0.716746	Akaike info criterion		2.237555
Sum squared resid	21.06271	Schwarz criterion		2.359204
Log likelihood	-46.22621	Hannan-Quinn criter.		2.282668
Durbin-Watson stat	0.898824			

This is probably as far as we can push the least-squares estimation of the monetary policy reaction function given by equation (7). In order to relax the assumption of both a constant inflation target and a constant natural real rate of interest, we need to turn to a time-varying parameter model. One of the most convenient approaches involves the state-space representation and the Kalman filter, to which we now turn.

3   An overview of the general state-space model[7]

The general form of a linear state-space model has quite a few components: unobservable states, observable data, shocks and system matrices. Despite this, it can always be written as a system of two equations. The first is a **measurement** equation (also known as the **signal** or **observation** equation):

$$y_t = H_t \beta_t + A_t Z_t + \varepsilon_t \tag{8}$$

where y_t is a ($n \times 1$) vector of observable variables, H_t is a ($n \times m$) matrix, Z_t is a ($n \times k$) matrix of exogenous variables, β_t is a ($m \times 1$) vector of (possibly) unobservable state variables, A_t is a ($k \times n$) matrix of parameters and ε_t is an observational error with E(ε_t) = 0 and var(ε_t) = R_t, where R_t is a known ($n \times n$) matrix.

The unobserved state variables β_t are generated by a first-order Markov process defined by the **transition** (or **state**) equation:

$$\beta_t = \mu_t + F_t \beta_{t-1} + S_t v_t \tag{9}$$

---

[7] For more details, interested readers are directed to Harvey (1989), Chapter 13 in Hamilton (1994a), Hamilton (1994b) and Koopman *et al.* (1999).

where F_t is a ($m \times m$) matrix, μ_t is a ($m \times 1$) vector, S_t is a ($m \times g$) matrix and v_t is a ($g \times 1$) vector of serially uncorrelated disturbances with E(v_t) = 0 and var(v_t) = Q_t, where Q_t is a known ($g \times g$) matrix.

The two disturbances ε_t and v_t may be contemporaneously correlated, meaning that:

$$\text{E}(\varepsilon_t v_t') = \Omega_t = \text{var}\begin{bmatrix} \varepsilon_t \\ v_t \end{bmatrix} = \begin{bmatrix} R_t & G_t \\ G_t' & Q_t \end{bmatrix} \tag{10}$$

In most cases, however, we have that the matrix of covariances $G_t = 0$, i.e., cov(ε_t, v_t) = 0.

The matrices H_t, R_t, F_t, S_t, Q_t and G_t are known as the system matrices. Most of the elements of these matrices will be fixed elements, mainly ones and zeros.

The disturbance ε_t in the measurement equation (8) is an error-in-measurement. The state-space form thus naturally lends itself to modelling systems with measurement error. In this case, the transition equation defines the signal β_t, which is unobservable and only measured with error (noise).

The model given by equations (8) and (9) is broader than often needed since many of the components will be either absent or at least will not be time-varying. For instance, the transition matrices F_t in equation (9) are usually time-invariant, as are the variances of the shocks Q_t. In addition, many models do not need to include any exogenous variables, Z_t.

As an aside, a useful concept in this area is the **signal-to-noise ratio**, λ, which is defined as $\lambda = \sigma_v^2/\sigma_\varepsilon^2$. The signal-to-noise ratio plays the key role in determining how observations should be weighted for prediction and signal extraction. In a large sample, filtering is equivalent to a simple exponentially weighted moving average; the higher is λ, the more past observations are discounted. When λ is zero, the level is constant and all observations have the same weight. We will return to the signal-to-noise ratio below.

There are two main benefits to representing a dynamic system in state-space form. First, the state-space form allows unobserved variables to be incorporate into, and estimated along with, the observable model. Second, state-space models can be analysed using the powerful recursive Kalman (Bucy) filter outlined next.

3.1 The Kalman filter

The Kalman filter due to Kalman (1960, 1963) and Kalman and Bucy (1961) is an algorithm for generating minimum mean square error forecasts in a state-space model. In particular, it is a recursive algorithm for sequentially updating the one-step ahead estimate of the state mean and variance given new information. Since the state-space model is a very general formulation for linear models in which time-varying parameters, measurement errors and missing observations can all be dealt with very easily, the Kalman filter is a very useful tool. As a by-product, if Gaussian errors are assumed, the filter allows the computation of the log-likelihood function of the state-space model. This, in turn, allows the model parameters to be easily estimated by maximum-likelihood methods.

4 Allowing for a time-varying inflation target: the state-space model

This exercise relaxes the assumption of a constant inflation target. Based upon equation (7), the time-varying parameter model to be estimated is given by:

$$i_t = (1 - \delta)(r^n + \pi_t^e + (\beta - 1)(\pi_t^e - \pi_t^*) + \gamma \tilde{y}_t) + \delta i_{t-1} + \varepsilon_t \tag{11}$$

$$\pi_t^* = \pi_{t\text{-}1}^* + v_t \tag{12}$$

$$\text{var}(v_t) = \lambda \, \text{var}(\varepsilon_t) \tag{13}$$

We note from equation (12) that the implicit inflation target, π_t^*, is assumed to follow a random walk. Such a specification represents a flexible way of capturing possible time-variation in the inflation target and is an established means of estimating other unobserved macroeconomic variables in the literature. By estimating the Taylor rule given by equation (11) in a state-space model, we allow the coefficients of the Taylor rule components to be estimated alongside the inflation target. As such, we relax the earlier assumption that the inflation target is constant over time by allowing it to vary according to a pre-determined random-walk process. Equation (13) links the variances of the two error terms, such that the variability of the implicit target is linked to the variability of the policy rate. The parameter λ thus becomes the signal-to-noise ratio.

Variation in π_t^* is governed by the signal-to-noise ratio parameter, $\lambda = \text{var}(v_t)/\text{var}(\varepsilon_t)$. If $\lambda = 0$, then the model collapses to the earlier specification with a constant inflation target. In this case, there is no link between the movements of the policy rate and the inflation target. We should note that the contribution of this variance to overall variability is likely to be very small. Consequently, the estimate of $\text{var}(v_t)$ is estimated very imprecisely and with downward bias. In turn, this means that the maximum-likelihood (ML) estimate of $\text{var}(v_t)$ is biased toward zero. This problem, referred to as the 'pile-up problem' by Stock (1994), implies that the estimate of the signal-to-noise ratio, $\lambda = \sigma_v^2/\sigma_\varepsilon^2$, is also biased towards zero.[8]

4.1   Creating and defining state-space models in EViews

To specify and estimate a state-space model in EViews, you have to create a state-space object first. Select **Objects**/**New Object.../SSpace** from the main toolbar or type `sspace` in the command window. EViews will create a state-space object and open an empty state-space specification window. This window will allow you to use keywords and text to describe what EViews calls the signal equation (and everyone else the observation or measurement equations), state (or transition) equations, error structures, initial conditions and, if desired, parameter starting values for estimation. I have also given a name to this specific state-space object, which is `ss_equation7_prior`.

The **SSpace** object offers specification and estimation options for single and multiple dynamic equations in state-space form and navigates the user to the application of the built-in Kalman filter. Exogenous variables can be included in the state-space equations and variances for all equations can be specified in terms of model parameters. Model output is available in tabulated format or in graphs. New series can be created from the model output and subsequently stored as new objects for further analysis. These typically include one-step ahead and smoothed states and signals, filtered states and corresponding residuals and errors.

The representation of the state-space model used for estimation as a **SSpace** object is as follows:

---

[8] Stock and Watson (1998) provide a procedure that overcomes this problem. Note that we did not apply this procedure in our exercise (yet).

```
@param c(1) 0.86 c(2) 1.59 c(3) 2.55 c(4) 0.51

@ename e1
@ename v1

@evar var(e1) = c(4)
@evar var(v1) = (1/40)*c(4)

@signal rate = (1 - c(1))*(3.25 + exp_reuters1_s + (c(3) -
1)*(exp_reuters1_s - sv1) + c(2)*gap100) + c(1)*rate(-1) + e1

@state sv1 = sv1(-1) + v1

@mprior prior_pi
```

The **SSpace** object in EViews that we have just created consists of several parts. The first line informs EViews about the starting values of the parameters that are to be estimated.[9] Finding good starting values for coefficient estimates is important. There are – at least – three ways of going about finding starting values. The first is sheer luck. It may turn out that our current starting values may yield convergence of the maximum-likelihood estimates. Incidentally, note that it is – as always with maximum-likelihood estimation – good practice to start from a number of starting values to see if they yield the same result (although this may be a dangerous – and frequently depressing – strategy!). The second approach involves experimenting with different starting values until you come across numbers that result in sensible results (data-mining?). Such a 'trial-and-error' approach can, of course, be time-consuming (and may not always be crowned with success). The final approach is to 'calibrate' error variances and coefficient estimates (data-mining again?). For the question at hand, for example, the starting values are based on the OLS estimates in Table 2. The least obvious starting value is probably the one for `c(4)`, which is based on the standard error of the regression in the bottom panel of Table 2. This is denoted `S.E. of regression` by EViews. Its value is equal to 0.716746 which, when squared, yields the starting value for the variance of `e1`, equal to 0.51.

In the second part (the next two lines), the error terms for the measurement and state equation (`e1` and `v1`) are named using the keyword `@ename`. Contrary to other objects in EViews, error terms are not included in the equation of a state-space object unless they have been explicitly specified. In this case, we include an error term for both the signal and the state equation.

The third part (the following two lines) specifies the variances. Error variances are preceded by the `@evar` keyword and may be constants or expressions in terms of unknown parameters. In our model, the variances are expressed as functions of the unknown coefficient `c(4)`. We can see that the variance of the observation or measurement equation is given by `c(4)`, while the variance of the transition equation is given by `(1/40)*c(4) = 0.025*c(4)`. This characterisation of the error variances of the transition equation mirrors their definition in equation (9), with $\lambda = 0.025$. In other words, for the initial

---

[9] Evaluation of the Kalman filter, smoother and forecasting procedures all require that we provide the initial one-step ahead predicted values for the states $\beta_{1|0}$ and variance matrix $P_{1|0}$. Unless stated otherwise, the starting values for the parameters are those specified in the coefficient vector c in the current EViews workfile. It is therefore always advisable to specify appropriate starting values.

estimation, Klein fixed λ, the signal-to-noise ratio, at $\lambda = 0.025$ (= 1/40). Using the `@evar` command, one can also specify covariances between error terms, for example, `@evar cov(e1,v1) = 0.5`. Variances and/or covariances that have not been specified are assumed to be equal to zero. If an error term is named, but there are no corresponding `var = ` or `@evar` specifications, the missing variance or covariance specifications will be set to the default values of zero. In many instances it is customary to express the variances as exponential functions of the unknown coefficients, where the exponential function is used to guarantee non-negative variance estimates (why?).

In the fourth part (the next two lines), the measurement and state equations are defined. The `@signal` keyword specifies the measurement equation for the dependent variable and includes an unobserved state variable, denoted `sv1`, to represent the (time-varying) inflation target as well as a disturbance, `e1`, which corresponds to ε_t and has been declared by the `@ename` command. The specification of the state equation should be obvious when we compare it to equation (11). If there is no error or error variance, the signal equation is assumed to be deterministic. Signal equations may include expressions of the dependent variable, but no current values or leads of the signal variables. Nonlinearities in the states and leads or lags of states are also not allowed. Note that the `@signal` keyword may be dropped. Finally, the 3.25 in the measurement equation corresponds to the constant natural real rate of interest.

The `@state` equation defines the random walk for the unobserved state variable, π_t^*, just as described in equation (12) State equations should not include expressions of unobserved components (like `log(sv1)`) nor lags or leads for signal equation dependent variables, but may contain (possibly non-linear) transformations of exogenous variables. They should be linear in the one-period lag of the states, where the one-period lag restriction is easily circumvented by including new state variables for higher-order lags. If there is no error or error variance, the state equation is assumed to be deterministic.

Error terms should necessarily be named before specifying the state and observation equation. But one can simply add the variance structure to define an error term, like in `@state sv1 = sv1(-1) + [var = (1/40)*c(4)]`. This is called the error variance specification, whereas the form shown above is the named error approach.[10] In other words, the state-space object could be equally well re-written as follows.

```
@param c(1) 0.86 c(2) 1.59 c(3) 2.55 c(4) 0.51

@signal rate = (1 - c(1))*(3.25 + exp_reuters1_s + (c(3) -
1)*(exp_reuters1_s - sv1) + c(2)*gap100) + c(1)*rate(-1) +
[var=c(4)]

@state sv1 = sv1(-1) + [var=(1/40)*c(4)]

@mprior prior_pi
```

---

[10] The main drawback of the error variance specification is that it does not permit correlation between errors in different equations. By default, EViews assumes that the covariance between different errors in different equations is zero. The 'named error' approach is useful when we require a more flexible variance-covariance structure.

By default, EViews will handle initial conditions for us. We may, however have prior information about the values of $\beta_{1|0}$ and $P_{1|0}$ (cf. Footnote 9). In this case, we can create a vector or matrix that contains the appropriate values, and use the `@mprior` or `@vprior` keywords to perform the assignment for the initial mean and variance values respectively. The final command therefore tells EViews that I have defined a (1×1) vector `prior_pi` which contains an initial value for π_t^*, equal to 5.09. In this exercise, we only have one unobserved state variable, `sv1`, representing the time-varying inflation target, so there is only one initial value to be specified. Note that in models with more than one unobserved state variable, each state variable will have to be initialised.

4.2 Estimation results

The first part of the analysis will proceed under the assumption that the natural real rate is constant at 3.25 per cent. Later sections of this exercise will relax this assumption and investigate the robustness of the results to this assumption. The estimation results, which are very close to those in Table 1 in Klein (2012, p. 7) are given in Table 3 below

**Table 3: Estimated response function of the
South African Reserve Bank (2001 Q1 – 2011 Q4)**

Sspace: SS_EQUATION7_PRIOR
Method: Maximum likelihood (Marquardt)
Sample: 2001Q1 2011Q4
Included observations: 44
User prior mean: PRIOR_PI
Convergence achieved after 25 iterations

	Coefficient	Std. Error	z-Statistic	Prob.
C(1)	0.800210	0.071914	11.12738	0.0000
C(2)	1.937002	0.667388	2.902361	0.0037
C(3)	3.921966	1.608288	2.438596	0.0147
C(4)	0.383625	0.092371	4.153070	0.0000
	Final State	Root MSE	z-Statistic	Prob.
SV1	6.264029	0.329956	18.98446	0.0000
Log likelihood	-50.57794	Akaike info criterion		2.480815
Parameters	4	Schwarz criterion		2.643014
Diffuse priors	1	Hannan-Quinn criter.		2.540967

The bulk of the output from state-space estimation in EViews should look familiar from other EViews estimation objects. The information at the top describes the basics of the estimation: the name of the **SSpace** object (`ss_equation7_prior`), estimation method (maximum likelihood using the Marquardt optimisation algorithm), the sample (2001 Q1 to 2011 Q4), the number of included observations in the sample (44), the prior mean (`prior_pi`), convergence information (EViews needed 25 iterations to achieve convergence towards a solution) and the coefficient estimates.

In terms of the estimation results, the estimated equation is consistent with common beliefs about the Taylor rule. All the estimated coefficients are positive (and statistically significant). The parameter β in equation (7), equivalent to `c(3)` in Table 3, is positive and significantly greater than 1, indicating not only that the Taylor principle holds, but also that

the SARB has been quite aggressive in responding to deviations of inflation expectations from its target rate over the sample period. The parameter γ in equation (7), equivalent to `c(2)` in Table 3, is also positive and significantly different from zero, suggesting that the SARB also considers output gap stabilisation as a secondary objective of monetary policy. The parameter δ in equation (7), which measures the degree of interest-rate smoothing, is equal to 0.800, indicating a relatively high degree of interest rate smoothing. Finally, the coefficient `c(4)` is the variance of the error term for the measurements and state equations. We therefore estimate that $\mathrm{var}(\varepsilon_t) = \sigma_\varepsilon^2 = 0.384$.

But Klein cautions that the estimation results should be treated with some care in light of the uncertainties regarding the 'true' values of the output gap and the natural interest rate.

The bottom part of the output view reports the maximised log-likelihood value, the number of estimated parameters and the associated information criteria.

On the other hand, some parts of the output are new. In particular, EViews displays the final one-step ahead values of the state vector, $\beta_{T+1|T}$, equal to 6.264, and the corresponding root mean square error values (square roots of the diagonal elements of $P_{T+1|T}$), equal to 0.330. In case you are interested in the entire path of the state (signal) vector and covariance matrix, EViews provides a variety of views and procedures for examining the state (signal) results in greater detail.

After estimation, EViews allows you to create new series based on the results of the estimation process. These can be generated using the appropriate keywords (in the command window) or they can be selected from the menu screens. As such, you generate the estimated state (and signal) series by clicking on **Proc** at the top of the **SSpace** object box and then selecting either **Make Signal Series** or **Make State Series**.
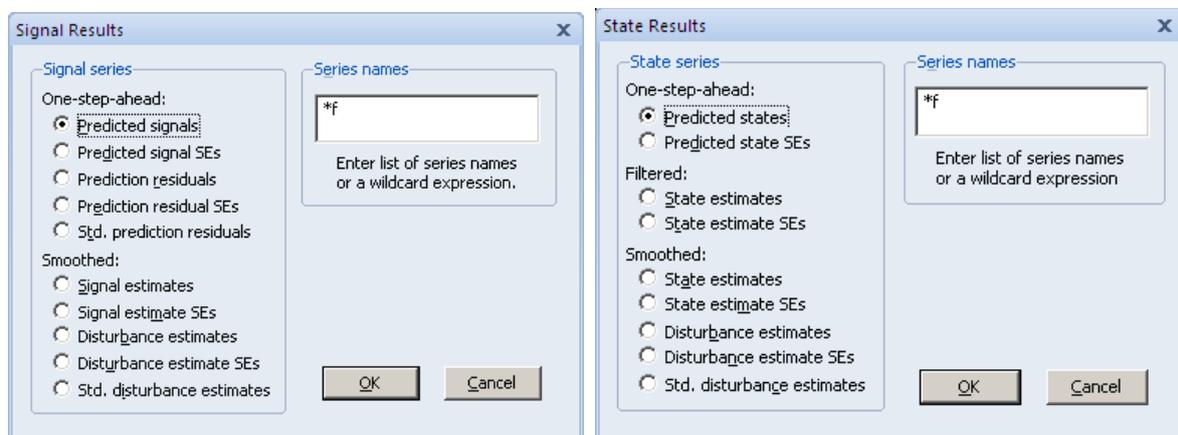
In our case, we would like to plot the filtered state series, `sv1`. For plotting, we click on **View** at the top of the **SSpace** object box and select **State Views** and **Graph State Series…** and **Filtered: State estimates**. The estimated implicit inflation target is given in Figure 2 below. I have added the upper and lower bands of the SARB's inflation range manually for comparison. This is done by right-clicking on the chart and selecting **Add lines & shading…** once for the upper horizontal (dashed) line at 6 and again for the lower horizontal (dashed) line at 3.

**Figure 2: Estimated implicit inflation target from the baseline model
(2001 Q1 – 2011 Q4)**

Filtered State SV1 Estimate
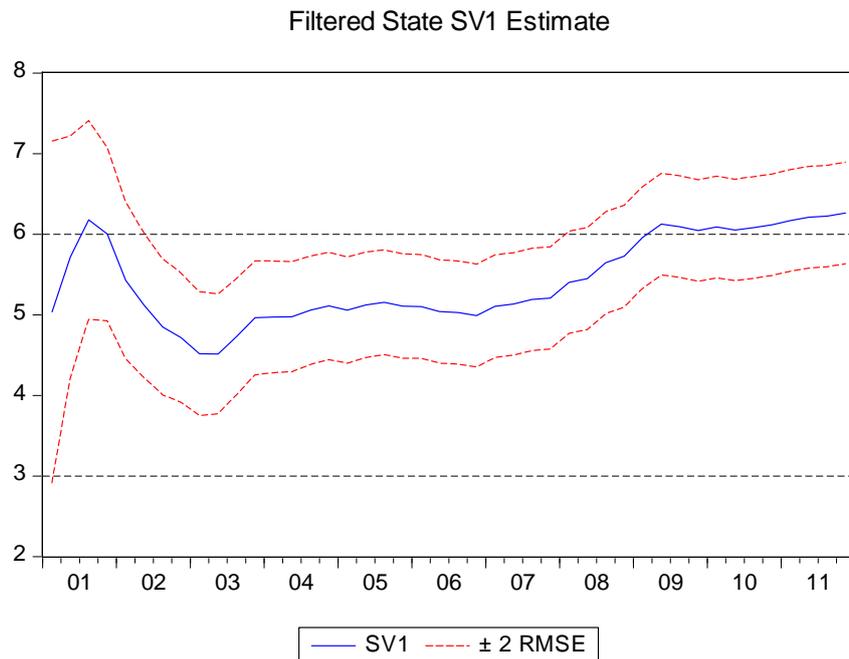
We can see that while the SARB's target range for inflation is defined as 3 to 6 per cent, the estimated implicit inflation target has generally been in the upper portion of that range since the introduction of IT. The results indicate that the implicit inflation target, at around 6 per cent, was high in 2001 and 2002, when the South African rand came under pressure on foreign-exchange markets. Still, the estimate moderated substantially over the course of 2003, falling to some 4.5 per cent. Over the period from 2004 to 2007, the implicit inflation target was more or less stable at around 5 per cent, but has been increasing since 2007 Q4 to around 5.5 per cent in 2008 Q3. With the onset of the financial crisis and the substantial monetary easing in late-2008 and early-2009, the estimated target recorded another increase, hovering close to – and maybe even above – the upper limit of the inflation target range at 6 per cent. This finding needs to be seen in light of the substantial variability of the output gap over the same period (Figure 1).

We save the filtered state estimate and its standard error by going to **Proc**/**Make State Series…** and then selecting, in turn, **Filtered: State estimates** and **Filtered: State estimate SE** and giving it an appropriate name (in my case, I have called them `sv1_base` and `sv1_base_se`).

A little manipulation of the underlying data as well as the chart options in EViews allows us to recover Figure 4 in Klein (2012, p. 7), reproduced as Figure 3 below. The band indicates the point estimates plus and minus one standard error. Before we can go and graph the series as an **Area Band**, we have to do a few more manual operations. In particular, we manually create the upper and lower confidence bands as:

$$sv1\_base\_upper = sv1\_base + sv1\_base\_se$$

and:

$$sv1\_base\_lower = sv1\_base - sv1\_base\_se$$

After we have created the two extra series, we open `sv1_base_upper`, `sv1_base_lower` and `sv1_base` in that order as a **Group**. You may now display an area band graph for the EViews **Group** object containing our three time series. EViews will construct bands from successive pairs of series in the group. If there is an odd number of series in the group, the final series will, by default, be plotted as a line. Select **View/Graph...** from the group menu, and then choose **Area Band** in the **Specific** listbox. The **Fill Areas** and **Lines & Symbols** dialogs under the **Graph Elements** group may be used to modify the characteristics of the lines and shades in your graph. After setting the appropriate options, the results should be the chart shown in Figure 3.

Figure 3: Estimated implicit inflation target from the baseline model (2001 Q1 – 2011 Q4)
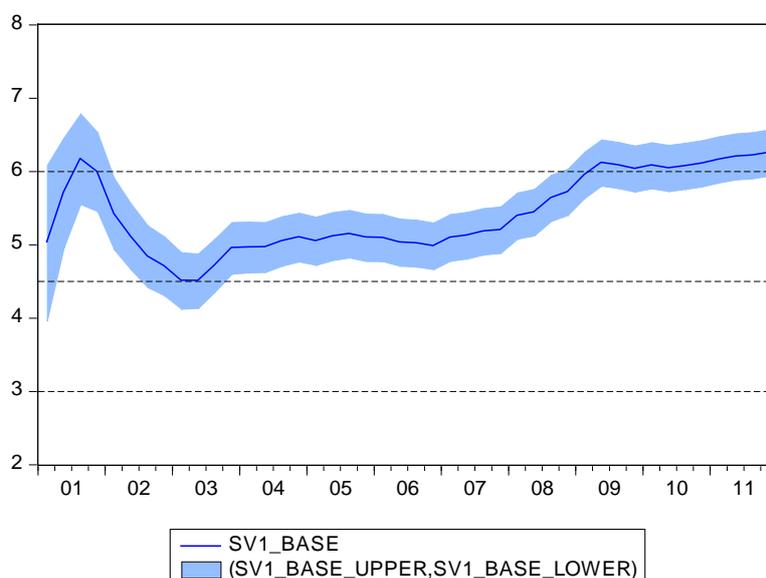
A Wald test rejects the null hypothesis that the inflation target is constant over time. This corresponds to the null hypothesis of H_0: var$(v_t) = 0$, which can be evaluated by testing whether $0.025*c(4) = 0$. Select **View/Wald Coefficient Tests...** and enter `c(4) = 0` in the dialog box (why?). The results of the test are shown in Table 4.

Table 4: Wald coefficient restriction test

Wald Test:
SSpace: SS_EQUATION7_PRIOR

Test Statistic	Value	df	Probability
Chi-square	17.24799	1	0.0000

Null Hypothesis: C(4)=0
Null Hypothesis Summary:

Normalized Restriction (= 0)	Value	Std. Err.
C(4)	0.383625	0.092371

Restrictions are linear in coefficients.

A Wald test of the restriction that the target is constant over time yields a χ^2-statistic of 17.248 (p-value: 0.000), meaning that we can reject the null hypothesis that var(v_t) = 0.

5   Assessing the sensitivity of the estimation results to pre-determined parameters

In our baseline estimations so far, three elements were pre-determined:

- the intercept of the Taylor rule equation, which is the assumed natural real rate of interest;
- the signal-to-noise ratio, λ, which determines the variability of the implicit inflation target relative to the variability of the policy rate; and
- the statistical process of the implicit inflation target, which was assumed to follow a random walk process

This section evaluates the degree to which deviations from these three assumptions affect the trajectory of the estimated implicit inflation target, starting with the statistical process of the implicit inflation target

5.1   Does the implicit inflation target follow a random walk or is it a stationary process?

At the moment, the statistical process of the implicit inflation target is assumed to follow a random-walk process, i.e., the coefficient on π^*_{t-1} in equation (8) is set to 1. But we can relax this assumption by reformulating equation (8) as:

$$\pi_t^* = \phi\pi^*_{t-1} + v_t \tag{14}$$

We can re-estimate the state-space model to assess whether the implicit inflation target follows a random walk or a stationary process and to what extent a change of this specification affects the trajectory of the implicit inflation target. The updated state-space model object is given by `ss_equation14`, which shows one change compared to the baseline specification. The state equation for `sv1` now includes an additional parameter called `c(5)` (with an initial value of 0.90):

```
@state sv1 = c(5)*sv1(-1) + v1
```

The results of this new state-space model are given in Table 5, which also includes the baseline results for comparison.

Table 5: Estimated response function of the SARB with a relaxation of the random-walk assumption (2001 Q1 – 2011 Q4)

Parameter	Baseline $\pi_t^* = \pi_{t-1}^* + v_t$	Alternative $\pi_t^* = \phi\pi_{t-1}^* + v_t$
δ	**0.800**	**0.846**
	(0.072)	(0.079)
γ	**1.937**	**2.542**
	(0.667)	(1.280)
β	**3.922**	3.132
	(1.608)	(1.790)
$\text{var}(\varepsilon_t)$	**0.384**	**0.393**
	(0.092)	(0.096)
$\text{var}(v_t)$	0.0096	0.0098
λ	0.025	0.025
ϕ		**0.999**
		(0.003)

Notes: Numbers in parentheses indicate standard errors. Numbers in the Baseline column replicate the results from Table 3. Coefficients that are statistically significant at the 5 per cent level of significance are denoted in boldface.

We find, on the basis of a Wald coefficient test, that the implicit inflation target does follow a random walk process, as the null hypothesis that ϕ (c(5)) = 1 cannot be rejected (p-value = 0.855). In contrast to Klein, however, I find that the rest of the estimated parameters change compared to the baseline specification. Consequently, the trajectory of the estimated implicit target (sv1_no_rw) is not very close to the baseline trajectory (Figure 4), especially at the beginning of the sample period.

Figure 4: Estimated implicit inflation target from the baseline and the random-walk models (2001 Q1 – 2011 Q4)
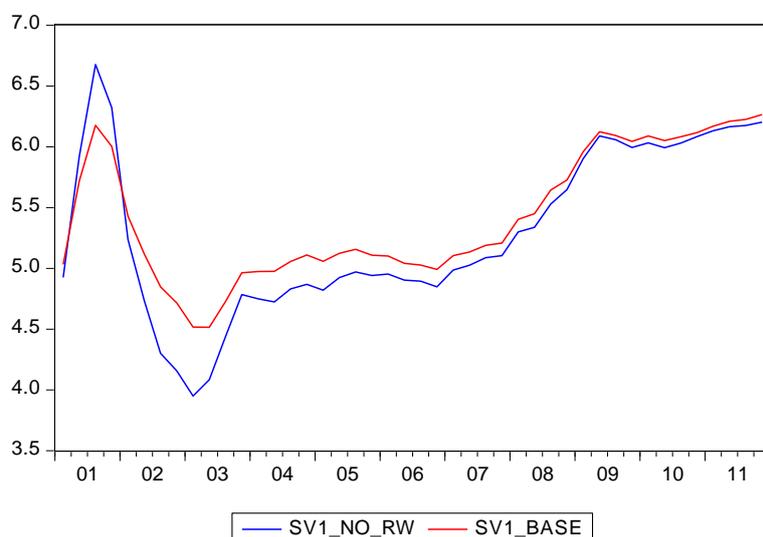
5.2 A different (constant) natural real rate of interest

To assess the sensitivity of the baseline results to the assumption of a constant natural real rate of interest equal to 3.25 per cent, which is the average real rate of interest in the sample, Klein considers two experiments. In the first, he considers a somewhat higher

(constant) natural real rate of interest equal to 4 per cent (75 basis points above the baseline real rate), while in the second he sets the (constant) natural real rate of interest equal to 2.5 per cent (75 basis points below the real rate).

Table 6: Estimated response function for the SARB with different natural real rates of interest (2001 Q1 – 2011 Q4)

	Baseline	Higher natural real rate of interest	Lower natural real rate of interest
Parameter	$r^n = 3.25$	$r^n = 4.00$	$r^n = 2.50$
δ	**0.800**	**0.800**	**0.800**
	(0.072)	(0.072)	(0.072)
γ	**1.937**	**1.937**	**1.937**
	(0.667)	(0.667)	(0.667)
β	**3.922**	**3.922**	**3.922**
	(1.608)	(1.608)	(1.608)
$\text{var}(\varepsilon_t)$	**0.384**	**0.384**	**0.384**
	(0.092)	(0.092)	(0.092)
$\text{var}(v_t)$	0.0096	0.0096	0.0096

Notes: Numbers in parentheses indicate standard errors. Numbers in the Baseline column replicate the results from Table 3. Coefficients that are statistically significant at the 5 per cent level of significance are denoted in boldface.

We find that even though the impact of different natural real rates of interest is insignificant in terms of the estimation results in Table 6 (up to three significant figures, at least), the impact on the estimated implicit inflation targets is still quite visible (Figure 5).

Figure 5: Estimated implicit inflation target with different levels of the natural real rate (2001 Q1 – 2011 Q4)
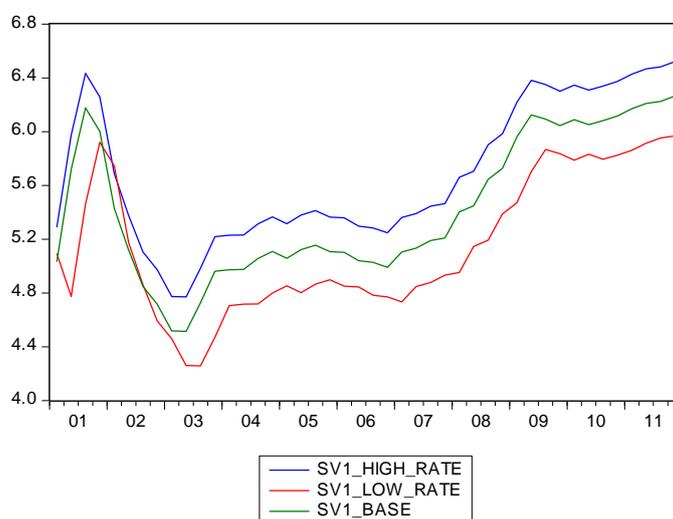
5.3 Alternative signal-to-noise ratios

The final sensitivity analysis evaluates different signal-to-noise ratios. In general, a higher (lower) λ implies that the implicit inflation target could change more (less) rapidly over the sample period. In addition to the baseline value of $\lambda = 0.025$, Klein considers two alternative values: $\lambda = 0.0125$ and $\lambda = 0.2$. The results in Table 7 allow us to investigate the

relative weights given to inflation and output stabilisation in the SARB's estimated reaction function. Both alternative models lead to a lower relative weight on inflation, as illustrated by the lower ratio of β compared to γ. Moreover, a higher λ is associated with a lower degree of interest-rate smoothing.

Table 7: Estimated response function for the SARB with
constant natural real rate and different signal-to-noise ratios (2001 Q1 – 2011 Q4)

Parameter	Baseline $\lambda = 0.025$	$\lambda = 0.0125$	$\lambda = 0.2$
δ	**0.800**	**0.833**	**0.742**
	(0.072)	(0.078)	(0.062)
γ	**1.937**	**2.192**	**1.719**
	(0.667)	(0.942)	(0.405)
β	**3.922**	**3.899**	**3.352**
	(1.608)	(1.926)	(0.991)
$\text{var}(\varepsilon_t)$	**0.384**	**0.419**	**0.277**
	(0.092)	(0.100)	(0.072)
$\text{var}(v_t)$	0.0096	0.0052	0.0554
β/γ	2.1	1.8	1.9

Notes: Numbers in parentheses indicate standard errors. Numbers in the Baseline column replicate the results from Table 3. Coefficients that are statistically significant at the 5 per cent level of significance are denoted in boldface.

The new implied trajectories for the implicit inflation target for the three signal-to-noise ratios are shown in Figure 6 below.

Figure 6: Estimated implicit inflation target with alternative signal-to-noise ratios
(2001 Q1 – 2011 Q4)
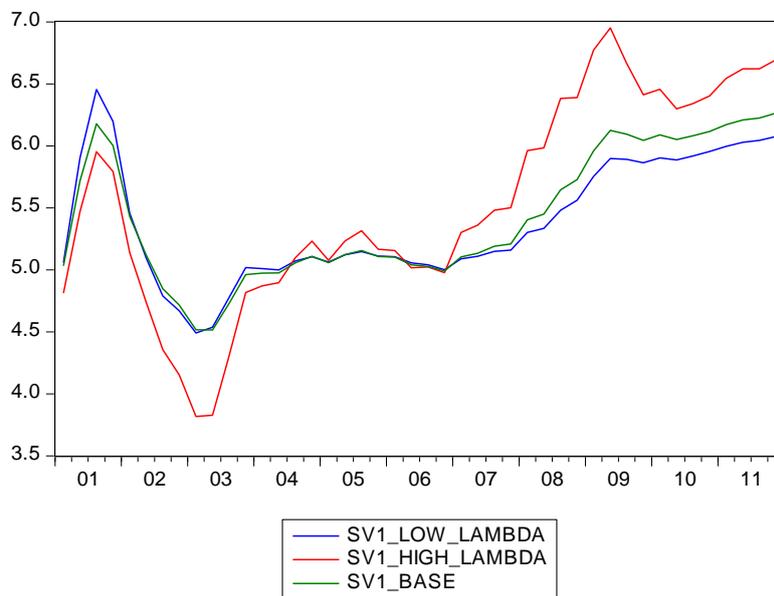
We can see that a change in the signal-to-noise ratio has a non-linear impact on the estimated trajectories. In particular, we can see that a higher λ is associated with a greater variability of the estimated target: the estimated target is below the baseline at the beginning of the sample, but above it at the end. In numbers, the high signal-to-noise ratio of $\lambda = 0.2$

results in the implicit inflation target increasing to a peak of 7 per cent in 2009 Q3 before moderating to around 6.6 per cent thereafter. The implicit inflation target is much more stable under the low signal-to-noise ratio of $\lambda = 0.0125$. Even so, there is a discernible upward drift in the estimated target towards the end of the sample period. Overall, the average estimated inflation target over the entire sample period from 2001 Q1 to 2011 Q4 is significantly lower under the low signal-to-noise ratio (6 per cent) than under the high signal-to-noise ratio (6.5 per cent).

5.4   Investigating a time-varying natural real rate of interest

The robustness analysis in Sections 5.1 through 5.3 assumed that the natural real rate of interest was fixed at different, albeit constant, levels over time. This assumption might be overly restrictive, particularly in light of the possibility that the natural real rate of interest has declined substantially since the outbreak of the global financial crisis. In general, we assume that the natural real rate of interest is closely linked to the potential growth rate of the economy (Laubach and Williams (2003)).

Klein therefore undertakes two further policy experiments. In the first, which he calls **partial variation** over time, the natural real rate of interest undergoes a one-off step-change downwards after 2008 Q4. He retains the value of 3.25 per cent for the natural real rate of interest from the start of the period until 2008 Q4, but then lowers the natural real rate to 2 per cent starting in 2009 Q1. This new natural real rate of interest variable is denoted `realrate1` in the workfile.

In a second experiment, he follows Laubach and Williams (2003) and estimates a state-space system in which the natural real rate of interest is linked to the estimated potential GDP growth. This model is referred to as **full variation** over time. He goes on to say that 'Under this approach the natural real rate depends on the trend growth of output, and is subjected to random disturbances (p. 11)'.

If I understand the model set-up correctly, the state-space model specification in EViews is as follows:

```
@param c(1) 0.86 c(2) 1.59 c(3) 2.55 c(4) 0.51 c(5) 0.5

@ename e1
@ename v1
@ename v2

@evar var(e1) = c(4)
@evar var(v1) = (1/40)*c(4)
@evar var(v2) = c(5)

@signal rate = c(1)*rate(-1) + (1 - c(1))*(c(2)*gap100 + (c(3)
- 1)*(exp_reuters1_s - sv1) + sv2 + exp_reuters1_s) + e1

@state sv1 = sv1(-1) + v1
@state sv2 = dhp_lnrgdp_long + v2

@mprior prior2
```

There are a few changes to the basic state-space model, which has been named `ss_equation7_full`. To begin with, I have made the natural real rate of interest in the signal equation time-varying by turning it into a state variable called `sv2`. As we now have a new state variable, we need to create a transition equation associated with it. Following Klein's description of the new dynamics of the natural real rate, the state equation is given by:

```
@state sv2 = dhp_lnrgdp_long + v2
```

such that the time-varying natural real rate of interest depends on the growth rate in potential output, `dhp_lnrgdp_long`, and a random disturbance term, whose error has a variance equal to `c(5)`. To account for the two state variables in the system, `prior2` is now a (2×1) vector with values given by 5.09 and 3.25.

One potential drawback with this set-up is that we relate the natural real rate of interest with the growth in potential output, which is very smooth. Indeed, the resulting estimated time-varying natural real rate over the period from 2001 to 2011 is probably too smooth to represent the natural real rate of interest (Figure 7). Other functional forms are, of course, possible.
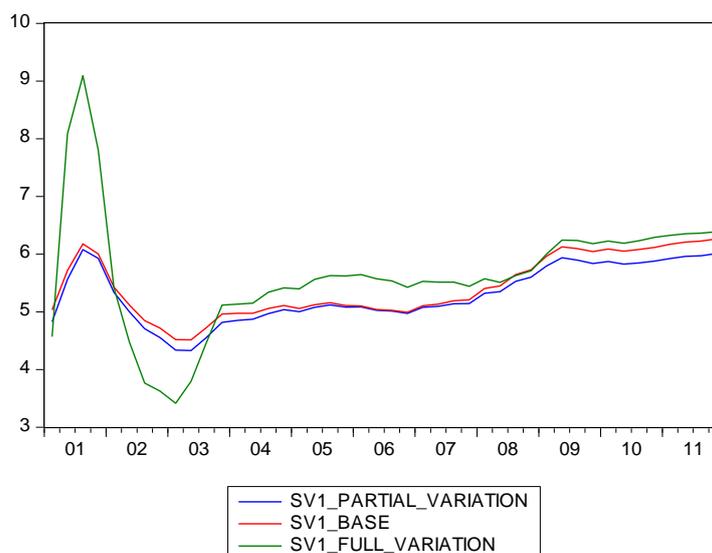
Figure 7: Estimated time-varying natural real rate of interest (2001 Q1 – 2011 Q4)
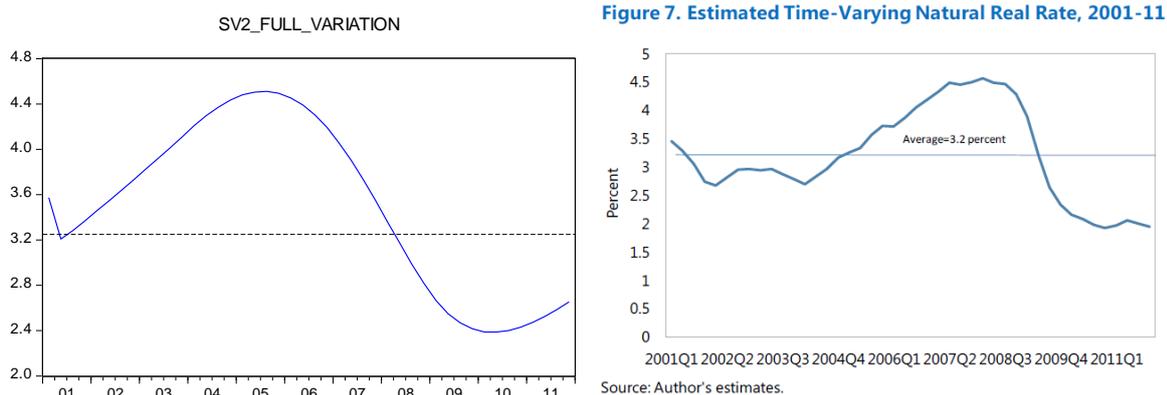
In comparison with Klein's Figure 7, the above chart is far too smooth, although the general trend is (kind of) the same. Moreover, the interest rate cycle peaks too soon: Klein finds that the natural real rate of interest increased quite sharply during the upturn in economic activity in the period from 2006 to 2008, and then declined significantly with the outbreak of the global financial crisis. In contrast, the natural real rate of interest in Figure 7 peaks as early as 2005. On a positive note, though, the magnitudes are comparable, in the sense that Klein's estimated natural real rate of interest series peaks at 4.5 per cent before falling to a low of 2 per cent.

For completeness, Figure 8 shows the three different implicit inflation targets.

Figure 8: Estimated implicit inflation target with time-varying natural real rate of interest (2001 Q1 – 2011 Q4)

The estimation results of these two experiments, as well as the baseline results, can be found in Table 8 below.

Table 8: Estimated response function for the SARB with time-varying natural real rate (2001 Q1 – 2011 Q4)

| Parameter	Baseline $r^n = 3.25$	Partial variation	Full variation
δ	**0.800**	**0.774**	**0.854**
	(0.072)	(0.059)	(0.070)
γ	**1.937**	**1.581**	1.765
	(0.667)	(0.494)	(1.056)
β	**3.922**	**3.441**	2.612
	(1.608)	(1.118)	(1.645)
$\text{var}(\varepsilon_t)$	**0.384**	**0.358**	0.501
	(0.092)	(0.087)	(5.051)
$\text{var}(\nu_t)$	0.0096	0.0089	0.0125
β/λ	2.0	2.2	1.5

Notes: Numbers in parentheses indicate standard errors. Numbers in the Baseline column replicate the results from Table 3. Coefficients that are statistically significant at the 5 per cent level of significance are denoted in boldface.

5.5   Forecasting exercise: actual versus predicted monetary policy rate

In order to complement the battery of robustness tests, Klein performs a dynamic forecast for the monetary policy rate to assess the extent to which the estimated state-space models describe the SARB's actual reaction function. In particular, he uses the state-space specifications of the baseline with a constant natural real rate of interest and the partial- as well as the full-variation specifications of the natural real rate of interest. The estimated state-space models are used to produce forecasts from 2005 Q1 until the end of the sample period.[11] This should be long enough to allow for sufficient cycles in the policy rate.

---

[11] It appears that Klein uses the models estimated up to 2011 Q4 to produce forecasts starting in 2005 Q1.

Forecasting using state-space models is done by clicking on the **Forecast** button. This opens up the **Forecast** dialog box, which is reproduced below.
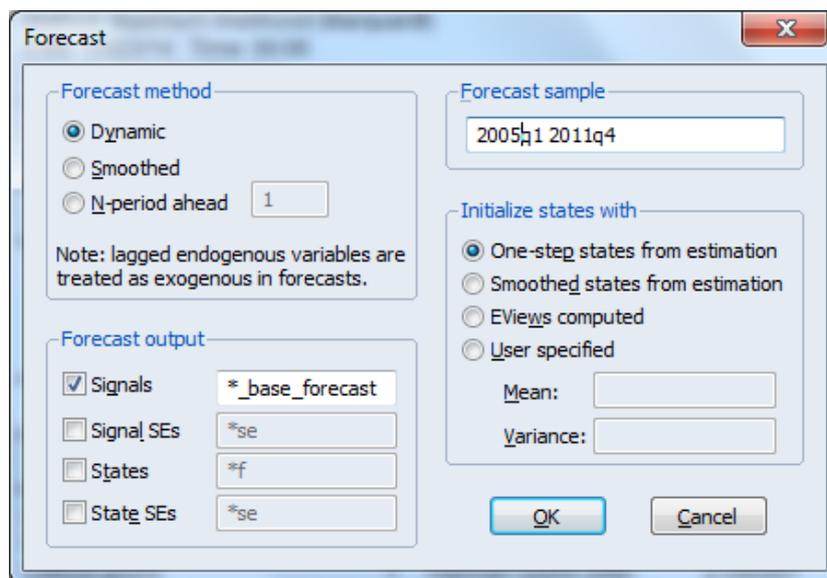
Note that what we are forecasting is not the state variable, but the signal, i.e., the actually observed time series of the base rate. I have given my forecast an appropriate name (`*_base_forecast`), such that the resulting EViews series in the workfile will be called `rate_base_forecast`. We can leave the other settings at EViews' defaults. Following Klein, we are performing a dynamic forecast over the period from 2005 Q1 to 2011 Q4, although the model is estimated on data that includes the forecast period. We then repeat the forecasting exercise using the state-space model allowing for partial variation of the natural real rate of interest as well as the full variation version of the state-space model. The actual monetary policy rate of the SARB as well as the predicted policy rates are shown in Figure 9.

Figure 9: Actual and predicted policy rates (2001 Q1 – 2011 Q4)
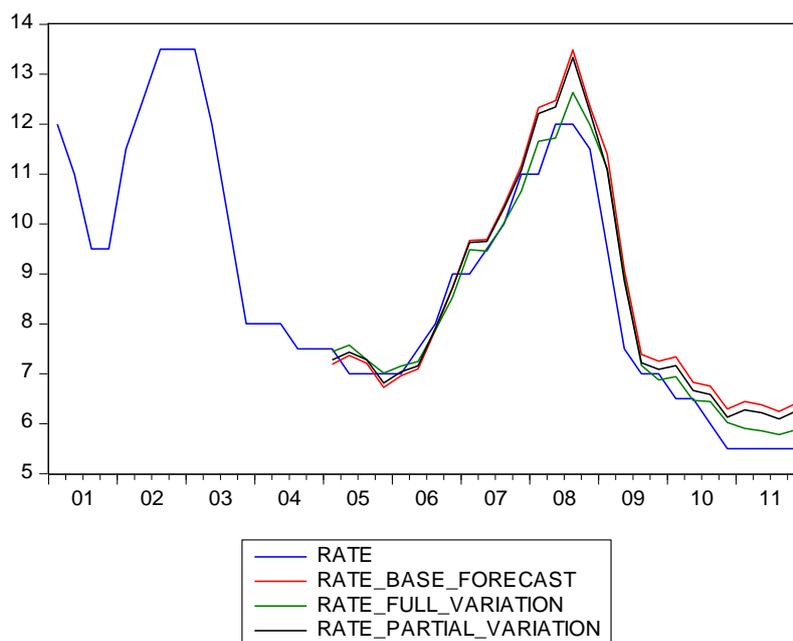
As discussed in Klein, the results show that the fitted rates follows the actual monetary policy rate quite well, even 'predicting' both the monetary tightening of 2006 to 2008 and the monetary relaxation after 2009. I should reiterate that this is not forecasting in the traditional sense and that this exercise is not intended for forecast use. At the same time, we should not that true out-of-sample forecasting with estimated state-space models can be very easily done.

6   Summary

This exercise illustrated an approach to estimating the implicit inflation target of a central bank and how it varies over time using state-space models and the Kalman filter. In particular, we estimated a simple Taylor rule in a state-space representation, where the coefficients of the Taylor rule components were estimated alongside the time-varying inflation target.

References and further reading

Andrews, D W K (1993), 'Tests for parameter instability and structural change with unknown change point', *Econometrica*, Vol. 61, No. 4, pages 821-56.

Andrews, D W K and Ploberger, W (1994), 'Optimal tests when a nuisance parameter is present only under the alternative', *Econometrica*, Vol. 62, No. 6, pages 1383-1414.

Clarida, R, Galí, J and Gertler, M (1998), 'Monetary policy rules in practice: some international evidence', *European Economic Review*, Vol. 42, No. 6, pages 1033-67.

Hamilton, J D (1994a), *Time series analysis*, Princeton, Princeton University Press.

Hamilton, J D (1994b), 'State-space models', Chapter 50 in Engle, R F and McFadden, D L (eds), *Handbook of econometrics, volume IV*, Amsterdam, Elsevier, pages 3039-80.

Harvey, A C (1989), *Forecasting, structural time series models and the Kalman filter,* Cambridge, Cambridge University Press.

Hansen, B E (1997), 'Approximate asymptotic p values for structural change tests', *Journal of Business & Economic Statistics*, Vol. 15, No. 1, pages 60-7.

Kalman, R E (1960), 'A new approach to linear filtering and prediction problems', *Journal of Basic Engineering, Transactions ASME, Series D*, Vol. 82, pages 35-45.

Kalman, R E (1963), 'New methods in Wiener filtering theory', in Bogdanoff, J L and Kozin, F (eds), *Proceedings of the first symposium on engineering applications of random function theory and probability*, New York, Wiley, pages 270-388.

Kalman, R E and Bucy, R S (1961), 'New results in linear filtering and prediction theory', *Journal of Basic Engineering, Transactions ASME, Series D*, Vol. 83, pages 95-108.

Klein, N (2012), 'Estimating the implicit inflation target of the South African Reserve Bank', *IMF Working Paper WP/12/177*. http://www.imf.org/external/pubs/ft/wp/2012/wp12177.pdf.

Koopman, S J, Shephard, N and Doornik, J A (1999), 'Statistical algorithms for models in state space using SsfPack 2.2', *Econometrics Journal*, Vol. 2, No. 1, pages 107-60.

Laubach, T and Williams, J C (2003), 'Measuring the natural rate of interest', *Review of Economics and Statistics*, Vol. 85, No. 4, pages 1,063-70.

Leigh, D (2008), 'Estimating the Federal Reserve's implicit inflation target: a state space approach', *Journal of Economic Dynamics & Control*, Vol. 32, No. 6, pages 2013-30. http://www.imf.org/external/pubs/ft/wp/2005/wp0577.pdf.

Lubik, T A and Schorfheide, F (2007), 'Do central banks respond to exchange rate movements? A structural investigation', *Journal of Monetary Economics*, Vol. 54, No. 4, pages 1,069-87.

Mise, E, Kim, T-H and Newbold, P (2005), ' On supoptimality of the Hodrick-Prescott filter at time series endpoints', *Journal of Macroeconomics*, Vol. 27, No. 1, pages 53-67.

Stock, J H and Watson, M W (1998), 'Median unbiased estimation of coefficient variance in a time-varying parameter model', *Journal of the American Statistical Association*, Vol. 93, No. 441, pages 349-58.