Chapter 7

7 Dynamic time inconsistency and the SARB*

7.1 Introduction

Realizing that South Africa has been targeting inflation since the February of 2000,[64] this chapter attempts to analyze whether the adoption of an inflation targeting regime has improved the time consistency of monetary policy in South Africa in terms of mean levels of inflation in the post-targeting period. Specifically, we try and deduce whether the South African Reserve Bank (SARB) could have produced lower average levels of inflation during the period of 2001:01 to 2008:02 if it had continued to pursue a monetary policy approach that it followed prior to 2000.

To do this, we would first need to obtain a framework that is in line with design of monetary policy under no precommittment to a rule. In this regard, we rely on the theory of dynamic time inconsistency. And then to get to the main question of ~~this chapter~~, understandably we first need to show that the SARB's monetary policy decisions were in line with a time inconsistent framework over the pre-inflation targeting period of 1960:01 to 1999:04. Econometrically speaking, this can be done by deriving restrictions imposed by the Barro and Gordon (1983) model of dynamic time inconsistency on a bivariate time-series model of Consumer Price Index (CPI) inflation and real Gross Domestic Product (GDP), and then testing these restrictions, both short- and long-run, based on quarterly data for South Africa covering the period of 1960:01 through 1999:04. And then, as far as answering the question posed above is concerned, this can done by forecasting inflation one-step-ahead over the period of 2001:01 to current, which in our case happens to be 2008:02.[65] And finally, checking, whether, on average, we would have obtained lower rates of inflation over the out-of-sample horizon of 2000:01 to 2007:04. However, it must be realized that to forecast out of the model, we need to ensure that the short-

64 See Sichei (2005) and Ground and Ludi (2006) for detailed reviews on the history of monetary policy in South Africa.

65 At the time, the paper was being written, data for the relevant variables were available till 2008:02 only.

*Published in South African Journal of Economics, (2010), 78(1):76-88.

and long-run restrictions obtained from the theory is consistent with the data, and, hence, the SARB was time inconsistent prior to targeting inflation.

The motivation for this analysis simply emanates from the need to evaluate the performance of the monetary authority during the inflation targeting period. Note, the issue is not whether the average level of inflation in the post inflation-targeting period is lower than the pre-inflation targeting period, but, more importantly, whether the inflation rate, on average, would have been higher or lower, if the monetary authority continued to behave in the way it had been prior to targeting inflation? The average level of CPI inflation over the period of 2000:01 to 2008:02 has been 6.17%, which is way lower than the average of 9.21% that prevailed in the pre inflation-targeting period of 1960:01 to 1999:04. Now a simple comparison of the levels of inflation in the pre and post inflation-targeting period would suggest that the inflation targeting regime has performed quite well in bringing down the average level of inflation in South Africa. However, this is an incorrect way of evaluating the performance of the new regime, because, ideally what we would want to deduce is whether the monetary authority could have done better or worse if it persisted with the policy structure of the old regime, into the period of 2000:01 to 2008:02. To the best of our knowledge, this is the first attempt to evaluate the performance of an inflation targeting monetary authority in this manner. Hence, an obvious extension of this chapter would be to carry out the analysis for other inflation targeting economies.

The remainder of the chapter is organized as follows: Section 7.2 lays the modified version of the Barro and Gordon (1983) model and derives the theoretical restrictions, while, Section 7.3 discusses the data and presents the empirical results and also carries out the forecasting exercise. Finally, Section 7.4 concludes.

7.2 The Modified Barro-Gordon (1983) Model

Recall, in Barro and Gordon's (1983) model, a policymaker wants to increase output or reduce unemployment, but does not pre-commit to a monetary policy rule, and, hence, is tempted to increase the output beyond the natural rate by creating unanticipated inflation or deviating from its pre-announced inflation rate in an attempt to exploit the expectational Phillips curve. In other words, the policy maker tends to be time inconsistent. However, given that private agents in the model have rational expectations, they can recognize this behaviour of the government and adjust their decisions accordingly. Therefore, in equilibrium, output is not lower than it would otherwise have been, and yet the rate of inflation is inefficiently high. This section presents the modified version of Barro and Gordon's (1983) model of time inconsistent monetary policy as can be found in Ireland (1999). However, unlike Ireland (1999), due to the lack of quarterly data on unemployment for South Africa, we model the supply side using a traditional Lucas-type supply curve rather than an expectational Phillips curve. Specifically, Barro and Gordon's (1983) model is modified by allowing the natural rate of output to follow an autoregressive process that contains a unit root and by incorporating control errors for the rate of inflation. While, the first extension allows for the real GDP to be non-stationary, as will be seen below, the second modification ensures transitory deviations between the actual real GDP and the natural rate of output.

As in the standard Lucas supply-curve, the actual output y_t fluctuate around the natural rate y_t^n in response to deviations of the actual inflation rate π_t from the expected inflation rate π_t^e as follows:

$$y_t = y_t^n + \alpha(\pi_t - \pi_t^e); \alpha > 0. \tag{7.1}$$

The natural rate of output, in turn, is assumed to fluctuate over time in response to a real (supply) shock ε_t according to:

$$y_t^n - y_{t-1}^n = \lambda(y_{t-1}^n - y_{t-2}^n) + \varepsilon_t;\ -1 < \lambda < 1; \varepsilon_t \sim iid\ \ N(0, \sigma_\varepsilon^2) \tag{7.2}$$

Hence, the change in the natural rate is allowed to follow an AR (1) process. The monetary authority cannot commit to a policy rule, but at the at the beginning of each period t = 0,1,2…, after the private agents have formed their expectation of inflation, π_t^e, but prior to the realization of the supply-shock, ε_t, the monetary authority chooses a planned rate of inflation π_t^p.Actual inflation for period t is then determined as the sum of π_t^p and a control error η_t, such that:

$$\pi_t = \pi_t^p + \eta_t; \eta_t \sim N(0, \sigma_\eta^2) \tag{7.3}$$

Note η_t is assumed to have a covariance of $\sigma_{\varepsilon\eta}$ with ε_t.

The policy maker chooses π_t^p in order to minimize a loss function that imposes penalty on variations of output and inflation around target values ky_t^n and zero:

$$L_t = \frac{1}{2}(y_t - ky_t^n)^2 + \frac{b}{2}\pi_t^2;\quad b > 0. \tag{7.4}$$

with k >1 and b>0 so that the policymaker wishes to push the actual output over the natural rate.

Using (7.1) and (7.3), the monetary authority's problem can be re-written as:

$$\min_{\pi_t^p} E_{t-1}\left[\frac{1}{2}\left((1-k)\,y_t^n+\alpha\left(\pi_t^p-\pi_t^e+\eta_t\right)\right)^2+\frac{b}{2}\left(\pi_t^p+\eta_t\right)^2\right] \tag{7.5}$$

where $E_{t-1}(.)$ denotes the expectation at the beginning of period t or at the end of period t-1. The first order condition for the above problem is:

$$\alpha E_{t-1}[(k-1)y_t^n+\alpha(\pi_t^p+\eta_t-\pi_t^e)]=bE_{t-1}(\pi_t^p+\eta_t) \tag{7.6}$$

Private agents are assumed to know the true structure of the economy and also understand the monetary authority's time-inconsistency problem. In equilibrium, therefore $\pi_t^p=\pi_t^e$, i.e., they correctly anticipate the authority's actions. Using the equilibrium condition and the fact that $E_{t-1}\eta_t=0$, due to rational expectations, (7.6) simplifies to:

$$\pi_t^p=\pi_t^e=\alpha AE_{t-1}y_t^n;\, A=\frac{k-1}{b}>0 \tag{7.7}$$

Equation (7.7), as in Barro and Gordon (1983), indicates that the inflationary bias resulting from the monetary authority's inability to commit depends positively on the expected natural rate of output $E_{t-1}y_t^n$.

Combining equations (7.1), (7.3) and (7.7) imply that the control error for inflation causes the actual output to fluctuate around the natural rate in equilibrium, i.e.,

$$y_t=y_t^n+\alpha\eta_t \tag{7.8}$$

Using (7.2) and (7.8), and defining $\Delta y_{t-1}^n=y_{t-1}^n-y_{t-2}^n$, we have:

$$y_t=y_{t-1}^n+\lambda\Delta y_{t-1}^n+\varepsilon_t+\alpha\eta_t \tag{7.9}$$

Meanwhile, equations (7.2), (7.3) and (7.7) imply that:

$$\pi_t=\alpha Ay_{t-1}^n+\alpha A\lambda\Delta y_{t-1}^n+\eta_t \tag{7.10}$$

Equations (7.9) and (7.10) separately indicate that both output and inflation are non-stationary respectively, by having inherited the unit roots from the underlying process defining the evolution of the natural rate of output. Putting (7.9) and (7.10) together implies that:

$$\pi_t-\alpha Ay_t=(1-\alpha^2A)\eta_t-\alpha A\varepsilon_t \tag{7.11}$$

which, in turn, shows that the linear combination of inflation and output is stationary, i.e. π_t and y_t are cointegrated. So based on equation (7.11), the modified version of the Barro and Gordon's (1983) model implies that inflation and real GDP are non-stationary, but cointegrated. Statistical tests of the cointegration constraint, implied by (7.11), will determine whether the modified Barro and Gordon (1983) model can explain the long-run behaviour of inflation and output in South Africa.

Taking first difference of (7.8) and (7.11), and then replacing out the first-differenced value of $\Delta y_t \times \alpha A$ from the first-differenced version of (7.8) into the first-differenced version of equation (7.11) yields:

$$\Delta\pi_t = \alpha\lambda\Delta y_t^n + \eta_t - \eta_{t-1} - \alpha A\varepsilon_t + \alpha A\varepsilon_{t-1} \qquad (7.12)$$

where $\Delta\pi_t = \pi_t - \pi_{t-1}$ and $\Delta y_t^n = y_t^n - y_{t-1}^n$.

Equation (7.12) in turn can then be re-written using a state-space representation as follows:

$$\varepsilon_t = F\varepsilon_{t-1} + Q\upsilon_t \qquad (7.13)$$

$$z_t = H\varepsilon_t \qquad (7.14)$$

where,

$$\varepsilon_t = \begin{bmatrix} \Delta y_t^n \\ \varepsilon_t \\ \varepsilon_{t-1} \\ \eta_t \\ \eta_{t-1} \end{bmatrix}; F = \begin{bmatrix} \lambda & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}; Q = \begin{bmatrix} 1 & 0 \\ 1 & 0 \\ 0 & 0 \\ 0 & 1 \\ 0 & 0 \end{bmatrix};$$

$$H = \begin{bmatrix} \alpha A & -\alpha A & \alpha A & 1 & -1 \end{bmatrix}; z_t = [\Delta\pi_t]; \upsilon_t = \begin{bmatrix} \varepsilon_t \\ \eta_t \end{bmatrix},$$

$$with\ E(\upsilon_t\upsilon_t') = \begin{bmatrix} \sigma_\varepsilon^2 & \sigma_{\varepsilon\eta} \\ \sigma_{\varepsilon\eta} & \sigma_\eta^2 \end{bmatrix}.$$

Further note, using (7.2), equation (7.12), after some algebra, can be re-written as:

$$\Delta\pi_t = \lambda\Delta\pi_{t-1} + \eta_t - (1+\lambda)\eta_{t-1} + \lambda\eta_{t-2} + (1+\lambda)\alpha A\varepsilon_{t-1} - \alpha\lambda A\varepsilon_{t-2} \qquad (7.15)$$

The within equation restriction appearing in (7.15) implied by the ARMA(1,2) process for the change in the actual inflation rate, summarizes the constraints that the modified version of the Barro and Gordon (1983) model imposes on the short-run behaviour of inflation. As with the long-run relationship, a statistical test of these restrictions will determine whether the modified model explains the dynamics of inflation that can be found in the South African data. This essentially boils down to using a likelihood-ratio test statistic for establishing the acceptance or the rejection of the short-run restrictions implied by equation (7.15) in relation to an unrestricted version of the ARMA (1,2) model of $\Delta\pi_t$,which looks as follows:

$$\Delta\pi_t = \phi_1\Delta\pi_{t-1} + \phi_2\varepsilon_{t-1}^\pi + \phi_3\varepsilon_{t-2}^\pi + \phi_4\varepsilon_{t-1}^y + \phi_5\varepsilon_{t-2}^y + \varepsilon_t^\pi \qquad (7.16)$$

At this stage, it is important to point out that we have reduced the two-variable model to a single variable $\Delta\pi_t$. This allows us to lower the number of parameters for the unrestricted ARMA (1,2) from 16 to 8, and in the process, help us reduce the difficulty of finding initial parameter values, via grid search, involved in estimating state-space models. Besides, in this chapter, in any case, we are more interested in studying the behavior of inflation over the period of 2000:01 to 2008:02.

7.3 Data and Results

In this chapter, we use seasonally adjusted quarterly time series data for real GDP and CPI inflation over the period of 1960:01-2008:02, both of which were obtained from the Quarterly Bulletins of the SARB. Note the base year is 2000, and we transform the real GDP series into its logarithmic values. In this section, we first discuss the tests of the long-run constraint and then move on to verifying the validity of the short-run restrictions.

7.3.1 Testing the Long-Run Restrictions

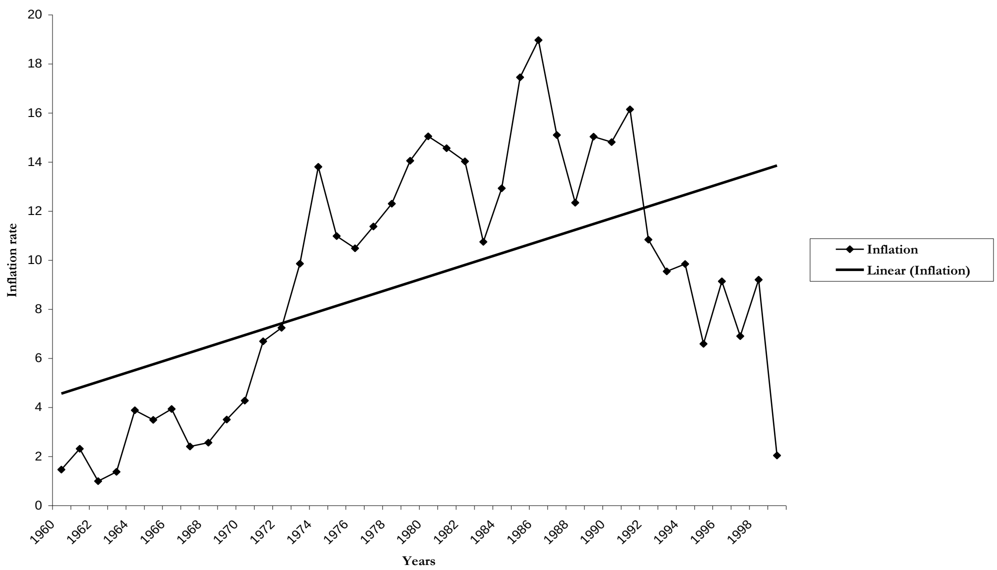

Figure 7-1: Linear trend of CPI inflation (2000=100).

Figure 7-2: Linear trend of LogRealGDP (2000=100).

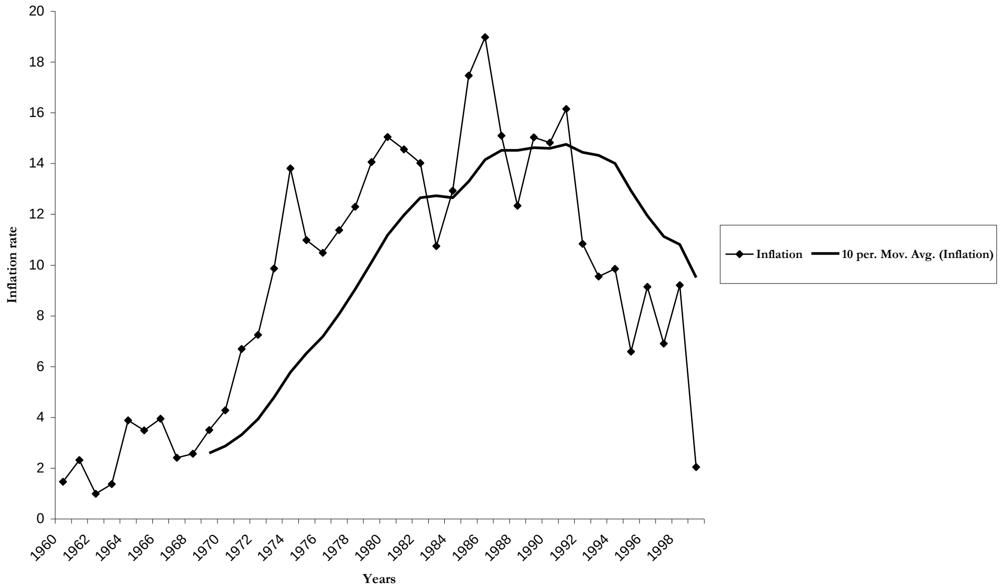

Figure 7-3: 10-year-centered Moving Average of CPI inflation (2000=100).

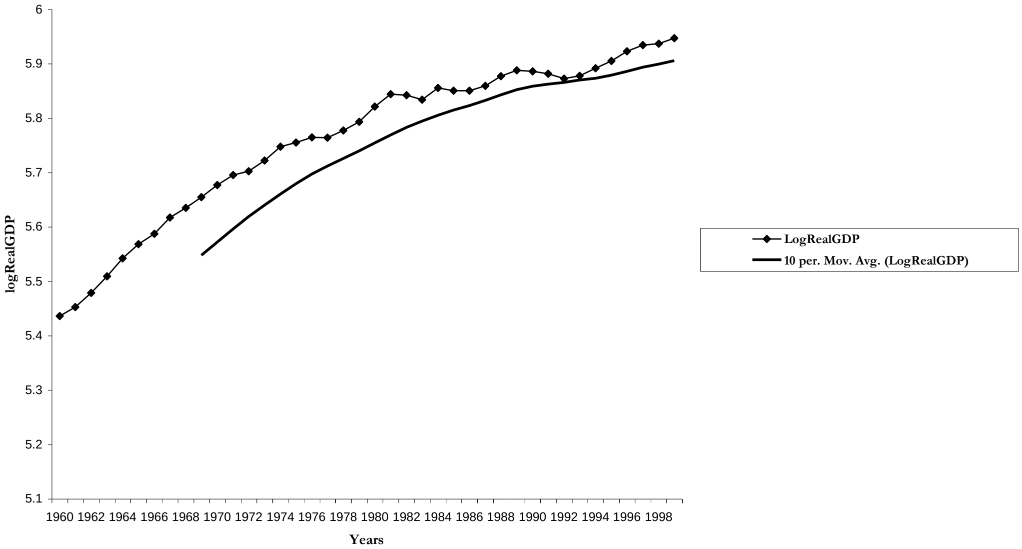

Figure 7-4: 10-year-centered Moving Average of LogRealGDP (2000=100).

Before we move to the formal tests of the long-run relationship, we present in Figures 1 through 4, the data plots for the inflation rate and the real GDP, and the associated trends based on a linear trend and a 10-period centred moving average over the period of 1960:01 to 1999:04. From the linear trends in Figures 7-1 and 7-2 for inflation and real GDP respectively, we find the variables to share a common positive trend. However, from Figures 7-3 and 7-4, based on the 10-period centred moving average, we find that the long-run inflation rate has experienced a downward movement since 1992 onwards, though output has continued to rise. Given this, it is likely that we might not find a cointegrating relationship between output and inflation over a shorter sample spanning the years of 1992 to 1999. But overall, for the whole period of 1960:01 to 1999:04, it is quite obvious that the two series are more than likely to be cointegrated.

Equations (7.9) and (7.10) indicate that the real GDP and the CPI inflation rate respectively, should be non-stationary. As can be seen from Table 7-1, based on the Augmented–Dickey–Fuller (ADF) test, the Dickey-Fuller test with GLS Detrending (DF-GLS), the Kwiatkowski, Phillips, Schmidt, and Shin (KPSS) test and the Phillips-Perron (PP) test, both the variables were found to follow an autoregressive process with a unit root, as the null hypothesis of a unit root could not be rejected for the variables, expressed in levels for the ADF, the DF-GLS and the PP tests, while for the KPSS test, the null of stationarity was rejected.

Table 7-1: Unit Root Tests.

Series	Model	ADF		PP	KPSS	DF-GLS	Conclusion
		$\tau_\tau \tau_\mu \tau$	$\phi_3 \phi_1$	$\tau_\tau \tau_\mu \tau$	$\tau_t \ \tau_\mu$	$\tau_t \ \tau_\mu$	
CPI		-0.81	13.69	-5.22***	0.37	-1.80	Non- stationary
	τ_μ	-2.19	23.13**	-4.29***	0.77	-0.90	
	τ	-0.68		-2.01**			
Δ CPI	τ_t	-10.59***	97.47***	-51.93***	0.12*	-13.95***	Stationary
	τ_μ	-10.42***	120.01***	-38.49***	0.31***	-1.41	
	τ	-10.45***		-38.21***			
Y	τ_t	-2.30	11.10**	-2.27	0.37	-0.41	Non- stationary
	τ_μ	-4.61***	21.23***	-4.73***	1.46	1.47	
	τ	4.08***		5.84***			
Δ Y	τ_t	-7.14***	59.98***	-12.97***	0.12***	-7.06***	Stationary
	τ_μ	-6.41***	82.09***	-12.33***	0.95	-3.54***	
	τ	-3.35***		-10.48***			

Note: ***(**)[*] indicates significance at 1 percent (5 percent) [10 percent] level.

To check for the cointegrating relationship between inflation and output implied by equation (7.11), we start of by using the Phillips-Ouliaris (1990) test. Results reported in Table 7-2 shows the estimate of γ, the coefficient from a regression of π_t on y_t, along with the statistics needed to test for a unit root in the residual from this regression. As can be seen, the hypothesis that inflation and output are not cointegrated can be rejected at the 1 percent significance level.

Table 7-2: Phillips-Ouliaris (1990) Cointegration Test (1960:01-1999:04).

γ	ρ	τ	Q	Z_{τ}
1.61	0.74	-4.92	0	-4.92***

Note: *** indicates significance at 1 percent level.

For the Phillips-Ouliaris (1990) test, Table 7-2 reports γ, the coefficient form the regression of inflation on real GDP, ρ, the coefficient from the regression of the residual from the inflation-output regression on its own lagged term and τ the conventional t-statistic for testing the hypothesis that ρ =1. Z_{τ} indicates the adjusted t-statistic by allowing for serial correlation in the regression error. The adjustment uses Newey and West's (1987) method to estimate the variance of the regression error and Andrew's (1991) method to select a value for the lag truncation parameter q required for the Newey and West (1987) estimator, assuming that the regression error is well approximated by a first-order autoregressive process.

One potential drawback of the residual-based Phillips-Ouliaris (1990) test is concerned with the fact that the sensitivity of the results might hinge on which variable (inflation or output), is used as the dependent variable in the initial regression. Here, however, equation (7.11) indicates that the hypothesized cointegrating relationship as suggested by the theoretical implications of the modified Barro and Gordon (1983) model is of the following form: $\pi_t - \gamma\ y_t$. This implies that we should be treating inflation as the dependant variable. Nevertheless, we check for the robustness of the results by using the Johansen (1988) test of cointegration, which, in turn, does not require a choice of normalization.

Based on a stable VAR[66] estimated with 5 lags[67], and allowing for no trend and intercept in the VAR or the cointegrating relationship, as suggested by the theory, we tested for cointegration using Johansen's (1988) maximum likelihood approach. Based on the Pantula Principle, both the Trace and the Maximum Eigen Value tests, showed that there is one stationary relationship (r = 1) between inflation and output at 1 percent level of significance. The results have been reported in Table 3, and they confirm the finding of the Phillips-Ouliaris (1990) test. The corresponding

[66] Stability, as usual, implied that no roots were found to lie outside the unit circle.

[67] The choice of 5 lags is based on the unanimity of the sequential modified LR test statistic, Akaike information criterion (AIC), and the final prediction error (FPE) criterion.

cointegrating vector relating inflation and real GDP, obtained from the Johansen (1988) approach, is found to be: $0.13\pi_t - 0.40y_t$.

Table 7-3: Estimation and Determination of Cointegrating Rank.

Null hypothesis	Alternative Hypothesis	Test statistic	0.05 critical value	Prob. **
			Trace Statistic	
r=0	r=1	24.62365	12.32090	0.0003
r=1	r=2	1.912715	4.129906	0.1962
Trace test indicates no cointegrating eqn(s) at the 0.05 level * denotes rejection of the hypothesis at the 0.05 level **MacKinnon-Haug-Michelis (1999) p-values				
			Maximum Eigenvalue Statistic	
r=0	r=1	22.71093	11.22480	0.0003
r=1	r=2	1.912715	4.129906	0.1962
Max-eigenvalue test indicates 1 cointegrating eqn(s) at the 0.05 level * denotes rejection of the hypothesis at the 0.05 level **MacKinnon-Haug-Michelis (1999) p-values				

Thus, the results reported in Tables 7-1 through 7-3 strongly support the long-run behaviour of inflation and real GDP implied by the restrictions obtained from the modified Barro and Gordon (1983) model.

7.3.2 Testing the Short-Run Restrictions

Focusing now on the short-run implications of the behavior of the change in the inflation rate, imposed by the modified Barro and Gordon (1983) model, Table 7-4 presents the maximum likelihood estimates of the model's parameters, which, in turn, is obtained by mapping (7-12) into a state space form, implied by (7-13) and (7-14). The likelihood function of the state-space model is then evaluated using Kalman filter.

The estimate of α =1.2716 suggest that the Lucas-supply curve is quite flat. Burger and Marinkov (2006) also draws similar conclusion about the slope of the curve using a VECM. Although, the parameters k and b are not identified individually, the estimate of A=(k-1)/b exceeds unity. With k>1, the result suggests that b<1, implying that the SARB placed more weight on its goal of output than on inflation over the pre-inflation-targeting era of 1960:01 to

1999:04. Again, similar observations have been made by Gupta and Naraidoo (2008), while estimating interest rate rules for South Africa in periods before the SARB moved into an inflation targeting framework. As expected, the estimate of λ is positive, though is not significant, as is the standard deviation of the real shock. The standard deviation that for the control error is, however, significant at the one percent level. Finally, the negative and significant estimate of the covariance $\sigma_{\varepsilon\eta}$ indicates that a positive shock to the natural rate tend to coincide with a negative shock to inflation. The estimate, thus, supports, the idea that ε_t represents a real shock or a supply-side disturbance.

Table 7-4: Maximum Likelihood Estimates.

Parameter	Estimate	Standard Error
α	1.27***	0.17
A	1.27**	0.61
λ	0.33	0.43
σ_{ε}	0.37	0.26
σ_{η}	2.44***	0.41
$\sigma_{\varepsilon\eta}$	-0.91**	0.45

Notes: (i) ***(**) indicates significance at 1 percent (5 percent) level;

(ii) L^c =-562.03; L^u = - 557.60

The within-equation restrictions that appear in (7.12) can be tested by comparing its fit with an unconstrained ARMA (1, 2) model identified in (7.16). The constrained model has 6 parameters, while the unconstrained model has 8. Thus, the theory places 2 restrictions on the univariate time series model for the stationary variable $\Delta\pi_t$. So, if L^c and L^u respectively, denotes the maximized value of the log-likelihood function for the unconstrained and the constrained model, then the likelihood ratio statistic LR= $2(L^u - L^c)$ has a chi-square statistic with 2 degrees of freedom under the null hypothesis that the constraints of the ARMA(1,2) model for $\Delta\pi_t$ holds. The LR statistic is: 2(562.03-557.60) = 8.86 < 9.21 (the 99 percent critical value for a χ^2 with 2 degrees of freedom). We can, thus, conclude that the model's short-run implications, as imposed by the theory, cannot be rejected at the 1 percent level of significance. This means that the data provides weak evidence of the theory in the short-run.

7.3.3 Evaluating the Inflation-Targeting Regime (2000:01-2008:02)

In this subsection, we evaluate the inflation targeting regime by trying to deduce whether the monetary authority could have done better or worse if it stayed time inconsistent over the period of 2000:01 to 2008:02 as well. To do this, we forecast the rate of inflation recursively using (7.12), first, based on new data generated from the one-period-ahead forecasts, and then second, based on the actual inflation rate that prevailed over this period. Ideally, because we are comparing across regimes, we would want to rely more on the forecasts generated from the forecasted values, rather than the original values. This is simply because, we are trying to analyze how the policymaker would have performed if it stayed time inconsistent, and hence, would not want to use the actual data that corresponds to the behaviour of the inflation rate in a different regime. However, just for the sake of completeness and comparison, we also forecast using the actual rate of CPI inflation.

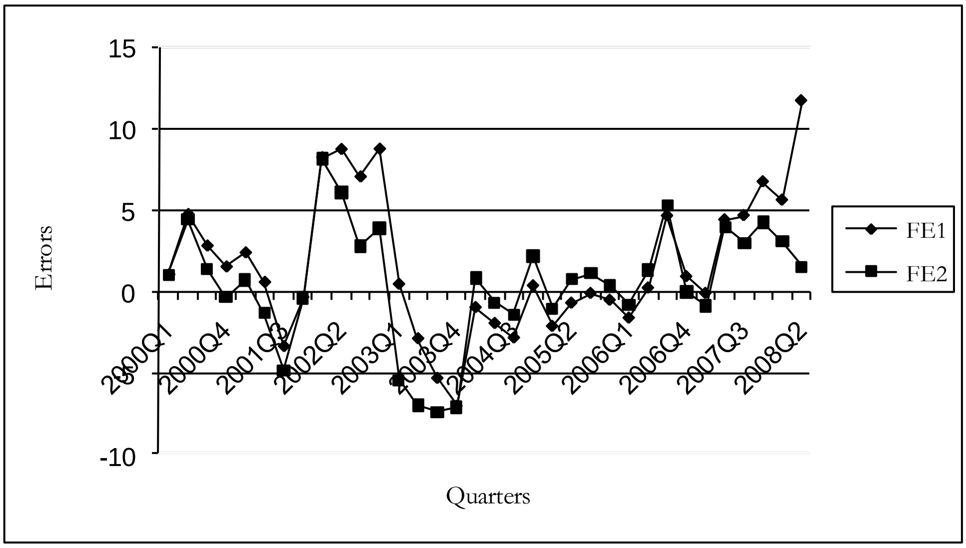

Figure 7-5: Differences between Actual and Forecasted Inflation.

Note: FE1(FE2) implies Forecast Errors based on forecast (actual) inflation.

Figure 7-5 plots FE1 and FE2, the one-step-ahead forecast errors based on forecasted values and actual values of the rate of inflation, respectively. We observe that the pattern of the movement of FE1 and FE2 are quite similar. Based on our calculations, the average value of the forecast errors based on the forecasted values of the inflation rate is 1.65 compared to 0.53 of the same when we use actual values. More importantly this implies, that if the SARB had continued to be time inconsistent, it could have produced on, average, an inflation rate which would have been lower by 1.65 % or 0.53% from what has prevailed over 2000:01 to 2008:02. Clearly then, the economy has been worse off in terms of the average levels of inflation experienced.

But then the big question is, whether these lower average rates of inflation, that could have been witnessed, are significantly different from what has actually been observed in the data on average? For this purpose, we resort to the Mincer and Zarnowitz (1969) regression. Note, for forecasts to be considered 'good' in an absolute sense they should not systematically under- or overpredict the rate of inflation. Formally, for a one-step-ahead forecast, we must have the following relationship:

$$E_{t-1}(\pi_t - \pi_{ft}) = 0 \tag{7.16}$$

where π_t and π_{ft} are the actual and the one-step-ahead forecasted values of the rate of inflation. Econometrically speaking, the Mincer and Zarnowitz (1969) regression boils down to testing the unbiasedness of the forecasts by regressing the forecasted values on the actual values of the variable under consideration.

Given the regression:

$$\pi_t = \beta_0 + \beta_1 \pi_{ft} + \nu_t \tag{7.17}$$

we test the joint unbiased hypothesis of: $\beta_0 = 0$ and $\beta_1 = 1$. At this stage it must be pointed out that the move to an inflation targeting regime can be considered to have improved the time consistency of monetary policy, then β_0 and β_1 in the Mincer and Zarnowitz (1969) regression should yield respectively, statistically significant values that are lower than zero and one. In other words, we would have $\pi_t < \pi_{ft}$.

Due to problems of possible serial correlation in the estimation of (7.17), Newey and West (1987) Heteroscedasticity and Autocorrelation Consistent (HAC) standard errors were calculated. Based on a Wald test, the probabilities of unbiasedness was found to be 0.37 and 0.46, depending on whether we use forecasted values or actual values to compute the forecasts.[68] This implies that the restrictions cannot be rejected and, hence, the model produces unbiased forecasts. However, more importantly this implies that, on average, the actual and forecasted values are not statistically different. Or in other words, the possible lower average rates of

[68] In an attempt to ensure that both π_t and π_{ft} were stationary, the regressions were re-estimated based on the first difference of the actual and forecasted values of inflation. The probabilities of unbiasedness based on the Wald test were found to be 0.81 and 0.16 respectively, based on forecasted values and actual values, implying that the restrictions, and, hence, unbiasedness cannot be rejected.

inflation that we could have obtained from a time inconsistent SARB is not significantly different from what the authority has achieved under the inflation-targeting regime.[69]

7.4 Conclusions

This chapter derives the econometric restrictions imposed by a modified Barro and Gordon (1983) model of dynamic time inconsistency on a bivariate time-series model of CPI inflation and real Gross Domestic Product (GDP), and tests these restrictions based on quarterly data for South Africa covering the period of 1960:01 through 1999:04. The results show that the data are consistent with the long-run implications of the theory of time-consistent monetary policy involving the two variables. However, as far as the short-run dynamics of the data is concerned, the evidence is weak. But importantly, when the model is used to forecast one-step-ahead inflation over the period of 2001:01 to 2008:02, i.e., the period covering the starting point of the inflation targeting regime till date, we, on average, produce lower rates of inflation, than those observed in the actual data.

However, based on the Mincer and Zarnowitz (1969) regression, we find that the possible lower average rates of inflation that we could have obtained from a time inconsistent SARB is not significantly different from what the authority has achieved under the inflation-targeting regime. But then again, realizing that in a general equilibrium framework, higher inflation rates can distort a host of other marginal decisions, besides the money demand, these so-called insignificant lower average rates of inflation could easily result in possibly quite large welfare losses. Given this, future research, should be aimed at quantifying the size of the welfare cost of inflation in a dynamic general equilibrium endogenous growth setting.

Significant or not, our results tend to show that retaining the pre-targeting monetary policy framework would on average have produced lower rate of inflation for South Africa, or in other words, inflation targeting seems to not have improved the time consistency of monetary policy expected out of such a regime. The results, thus, point to the fact that, perhaps the SARB needs to manage the inflation-targeting framework better than it has done so far. In this regard,

---

[69] All the estimations were carried out again for a a revised pre-targeting sub-sample of 1983:01-1999:04, with the starting date coinciding with that of Ortiz and Sturzenegger (2007). However, the changes in the parameter estimates were marginal and did not affect our general conclusions. These results are available upon request from the authors.

as pointed out generally by Demertzis and Viegi (2006, 2007, and 2008), a narrower, and possibly also a lower, target band could be of immense help in improving the central bank's credibility and causing inflation expectations to converge to a focal point, and hence, bring down the rate of inflation.

Chapter 8

8 Comparing South African Inflation Volatility across Monetary Policy Regimes: An Application of Saphe Cracking*

8.1 Introduction

Recent empirical evidence on the direct link of inflation targeting and particular measure(s) of economic performance, in our case inflation volatility, is at best mixed. While, Neumann and Hargen (2002), Petursson (2004), Vega and Winkelfried (2005) and Mishkin and Schmidt-Hebbel (2007) finds inflation targeting has led to low inflation volatility, studies such as Johnson (2002), Truman (2003), Ball and Sheridan (2005) and Fang et al. (2009) tends to suggest otherwise. Essentially, all these studies and other papers[70], analyzing the macroeconomic effects of inflation targeting, relies on empirical comparisons across inflation-targeting and non-targeting countries or within an inflation targeting country across monetary policy regimes, i.e., in the pre- and post-inflation targeting era. Though, comparisons across targeting and non-targeting countries are quite rational, studies that tend to compare within a country across regimes are flawed, and, hence, can only be viewed as providing preliminary evidence on the success or failure of inflation targeting. The reason is simple: When analyzing the effect of inflation targeting on the volatility of inflation, the real question is not whether the inflation volatility has increased or decreased since the central bank's decision to target inflation, but rather is the inflation rate more volatile than it would have been had the earlier monetary policy regime continued. To address this issue, we use what is called the cosine-squared cepstrum, and apply the technique to analyze the Consumer Price Index (CPI) inflation volatility in South Africa over the period of 1996:q1 to 2008:q3[71], which moved to an inflation targeting regime from the first quarter of 2000.[72]

[70] See Mishkin and Schmidt-Hebbel (2007) for a detailed survey.

[71] For the motivation behind the choice of the sample period, please refer to Section 3.

[72] For a detailed summary on the history of monetary policy in South Africa, please refer to Naraidoo and Gupta (2009).

*Forthcoming in Journal of Developing Areas.

Note in the mid- to late-1990s, the South African Reserve Bank (SARB) took a more eclectic approach to monetary policy, which, essentially involved monitoring a wide range of indicators, such as changes in bank credit extension, overall liquidity in the banking sector, the yield curve, changes in official foreign reserves, changes in the exchange rate of the rand, and inflation movements and expectations. This approach was enhanced in 1998 by the replacement of the discount window by the marginal lending facility of the repurchase system and consequently, the Bank rate was replaced by the repo rate. The SARB altered their previously eclectic approach in February of 2000, when the Minister of Finance announced that inflation targeting would be the SARB's sole objective. Currently, the Reserve Bank's main monetary policy objective is to maintain CPIX[73] inflation between the target-band of three to six percent, using discretionary changes in the repo rate as its main policy instrument.

Given that the SARB now targets inflation, if the move into the new regime or the abandonment of the older one affected inflation volatility, then observed inflation is represented as the sum of two series: (i) the series that would have eventuated if the SARB continued to pursue its more eclectic approach to monetary policy, and (ii) a second series associated with the direct impact of the regime change that arrived in late 1999 or early 2000. If the regime had any coherent effect, then the two series are likely to be well correlated. If the second series is found to be positively correlated with the series that would have eventuated, then it increases the volatility of inflation and is said to be in-phase. However, if the second series is negatively correlated or out of phase, volatility declines, since fluctuations in the series that would have eventuated are dampened. Generally speaking, if a time-varying stationary series is composed of two such series that are linear or well correlated, then the autocovariance function contains a global maximum at the zero lag and a local extremum at a lag corresponding to the date when the second series arrives. The problem with the autocovariance function is that it is difficult to detect a second series and determine the degree of its phase shift, given that a local extremum, in the case of an autocovariance function, appears as a broad cycle.

The cepstrum technique helps us to overcome such difficulties. Cepstra have been used successfully in detecting secondary influences in engineering, particularly communication theory.

[73] CPIX is defined as CPI excluding interest rates on mortgage bonds.

Intuitively, the cosine-squared cepstrum behaves like an autocovariance function, but with sharper resolution that helps in identifying the arrival and phase relationship of the secondary series with great precision. A local extremum appears as an impulse, instead of a broad cycle, the direction of which, in turn, determines whether the secondary series have increased or dampened volatility. Because of this, the cosine-squared cepstrum is well equipped to determine whether inflation volatility is greater than it would otherwise have been had the SARB continued to pursue its so-called eclectic approach to monetary policy decision-making.

The motivation to look into South Africa, specifically, emanates from the previous chapter which derives the econometric restrictions imposed by the Barro and Gordon (1983) model of dynamic time inconsistency on a bivariate time-series model of CPI inflation and real Gross Domestic Product (GDP), and tests these restrictions based on quarterly data for South Africa covering the period of 1960:q1 through to 1999:q4, i.e., for the pre-inflation targeting period. The results show that the data are consistent with the short- and long-run implications of the theory of time-consistent monetary policy. Moreover, when the model is used to forecast one-step-ahead inflation over the period of 2000:q1 to 2008:q2, i.e., the period covering the starting point of the inflation targeting regime till date we obtain, on average, obtain lower rates of inflation. Chapter 7 though is silent about inflation volatility that would have been generated over the inflation targeting era, if the SARB continued to be dynamically time inconsistent. But this is understandable, given that the Barro and Gordon (1983) model is essentially a framework for analyzing equilibrium inflation, and not its volatility. And this is exactly what this chapter tries to address, as far as South Africa is concerned. It must be realized that for a country seeking price stability, it is not only essential to obtain lower mean levels of inflation but also less volatility in inflation. Inflation volatility matters because high variability of inflation over time makes expectations about the future price level more uncertain, which, in a world with nominal contracts, induces risk premia for long-term arrangements, raises costs for hedging against inflation risks and leads to unanticipated redistribution of wealth. Thus, inflation volatility can impede growth, even if inflation on average remains restrained. Given this, this chapter is of paramount importance. To the best of our knowledge, this is the first attempt to study inflation volatility across monetary policy regimes by using the cosine-squared cepstrum.[74] Finally note the decision to use quarterly data rather than monthly data for the CPI inflation rates is essentially an

[74] The only other paper that uses, and in fact, introduced the cosine-squared cesptrum to economics is by Cunnigham and Vilasuso (1994). The authors used the technique to compare Gross National Product (GNP) volatility across exchange rate regimes for the US economy.

attempt to use the same data frequency as that of chapter 7 in drawing conclusions about the volatility of the inflation rates. For the same reason, we also used CPI rather than CPIX inflation. Besides, quarterly data on CPIX inflation is not available beyond 1997:q2 and our methodology requires data dating back to 1996:q1. For further details, refer to Section 3 below.

An extended version of this chapter would be to evaluate the performance of all the inflation targeting economies with respect to not only inflation volatility, but also the volatility of output and the monetary policy instrument used to maintain the target or the target-band. The rest of the chapter is organized as follows: Section 8.2 describes the data and our main findings, while, Section 8.3 concludes. Note, an appendix describing the technical details of a cepstrum has been presented at the end of the chapter.

8.2 Application to Inflation Volatility

If the move to the inflation targeting regime affected the volatility of the CPI inflation rate for South Africa, then the series is representable as the sum of the pre-2000 series and a secondary series arriving in late 1999 or early 2000. If the secondary series are perfectly in-phase with the pre-2000 series, then volatility increased relative to what it would have been under the eclectic monetary policy regime; if it is perfectly out-of-phase, then volatility decreased. Determining whether or not a secondary series arrived and it was in-phase or out-of-phase is a straightforward application of the cosine-squared cepstrum, discussed in the appendix of the chapter.

We begin by performing unit root tests on the CPI inflation rate, since the application of the cosine-squared cepstrum requires us to ensure that the data process is stationary. For this purpose we use the Augmented–Dickey–Fuller (ADF) test, the Dickey-Fuller test with GLS Detrending (DF-GLS), the Kwiatkowski, Phillips, Schmidt, and Shin (KPSS) test and the Phillips- Perron (PP) test. As can be seen from Table 1, seasonally adjusted (at annual rate) CPI inflation is found to follow an autoregressive process with a unit root, as the null hypothesis of a unit root could not be rejected for the inflation rate, expressed in levels for the ADF, the DF-GLS and the PP tests, while for the KPSS test, the null of stationarity was rejected. As the variable of interest is found to be non-stationary, we had to first difference the inflation rate series. Once the first-differenced CPI inflation series was generated and found to be stationary, we demean the series to avoid the dominance of the zero frequency components in the power spectrum, and then apply the five point Hanning-type cosine tapers to suppress possible sidelobes.

Table 8-1: Unit Root Tests (1996:Q1-2008:Q3).

Series	Model	ADF		PP	KPSS	DF-GLS	Conclusion
		$\alpha_t \alpha_\mu \alpha$	$\phi_3 \phi_1$	$\alpha_t \alpha_\mu \alpha$	$\alpha_t \alpha_\mu$	$\alpha_t \alpha_\mu$	
cpi_infl	α_t	-2.75	4.13	-2.89	0.13	-2.86	Non- stationary
	α_μ	-2.80	7.86***	-2.97**	0.13***	-2.72***	
	α	-1.24		-1.00			
D(cpi_infl)	α_t	-6.86***	23.53***	-8.02***	0.11***	-6.99***	Stationary
	α_μ	-6.86***	47.09***	-7.79***	0.23***	-6.75***	
	α	-6.92***		-7.86***			

Notes: (i) cpi_infl stands for CPI inflation;
(ii) D(cpi_infl) is the first difference of CPI inflation.

Next we use a Cooley–Tukey fast Fourier transform (FFT) and compute the natural logarithm of the sum of squares of the real and imaginary parts to form the log power spectrum. Note we start our analysis 15 quarters prior to the change in regime, so the first-differenced CPI inflation series begins at 1996:q2.[75] The choice of 15 quarters is decided based on the average distance between the troughs of the smoothed[76] log power spectrum, which was found to be approximately 0.0686 cycles per quarter, so that the lag associated with the secondary influence on the first-differenced inflation rate is approximately at a lag of 15 quarters, i.e., $\tau = 15$, arriving in the first quarter of 2000. Given that the seasonally (adjusted at annual rate) CPIX inflation rate was also non-stationary, using the first-differenced CPIX inflation rate over 1997:q3-2008:q3

---

75 The data for this chapter are from the quarterly bulletins available for download from the official website (http://www.reservebank.co.za/) of the SARB. All the computations are performed by using the Signal Processing Toolbox in MATLAB, Version R 2007a.

76 Following Bogert et al. (1963), the power spectrum is smoothed by applying a five-point centered moving average, since the moving average suppresses high frequency variations that would correspond to relatively long lags in secondaries.

revealed that the distance between the troughs of the smoothed log-power spectrum was approximately 0.06812 cycles per quarter on average, implying a τ = 15 as well. With CPIX inflation data only available from 1997:q2 and the regime change taking place in 2000:q1, we could not use the CPIX inflation data to obtain the cosine-squared cepstrum. Forming the inverse FFT (IFFT) and then the sum of squares again, and establishing the sign according to the sign of the real part, the cosine-squared cepstrum is formed, as discussed in Equation (A.6) in the appendix.

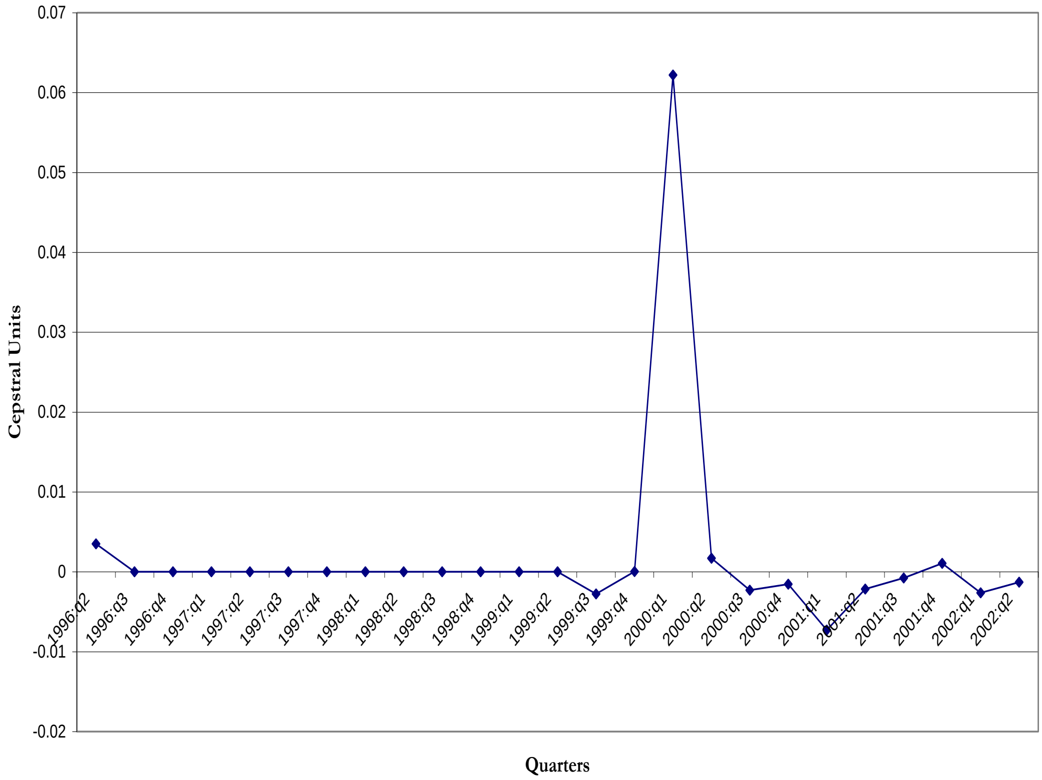

Figure 8-1: Shows the first 25 lags of the cosine-squared cepstrum for first-differenced CPI inflation.

For the sake of convenience, the time scale replaces the lag numbers. The first twelve points of the cepstrum for the first-differenced CPI inflation rate is rescaled to enhance the detail.[77] A delta function spikes prominently upward at the lag corresponding to the first quarter of 2000. This "arrival time" for the secondary series is consistent with the announcement of the Finance Minister that inflation targeting would be the SARB's sole objective. The unambiguous nature of the delta function and its positive sign provide evidence of a secondary influence on the South African CPI inflation series that not only matched it fluctuation by fluctuation, but was also exactly in-phases with the CPI inflation series. In other words, the secondary influence that arrived in early 2000 has increased the fluctuations/volatility in the CPI inflation rate.

Chapter 7 indicated that had the SARB continued to be time inconsistent it would have produced lower average levels of inflation in the post inflation targeting period. Here we show that the inflation-targeting regime has increased the volatility of inflation relative to what it would have been under the eclectic monetary policy regime. So in other words, it seems, at least until now, the SARB has failed to attain its primary objective of inflation targeting -- price stability in terms of both the mean and variance of inflation.

8.3 Conclusions

The positive delta function in the cosine-squared cepstrum has provided evidence that the inflation targeting regime in South Africa began to impact CPI inflation in the first quarter of 2000, making the same more volatile than it would have been had the SARB pursued the more eclectic monetary policy approach, which was in place prior to the 2000:Q1. The cosine-squared cepstrum, thus, shows that it can be successfully used in applications where the effects policy changes on time series data can be modelled as additive time series that are well correlated with the variable under study. However, there are limitations to this approach: First, the increased inflation volatility under the new monetary policy regime may not be permanent, but rather a pulse-like response in the inflation rate, and; Second, at times the methodology might require additional economic insight to isolate the possible economic causes of the event, without

[77] The first twelve cepstral estimates for the first-differenced CPI inflation rate is set equal to zero to enhance the detail. The above operation removes the delta function at the origin, which, in turn corresponds to the global maximum.

recovering the secondary series. Fortunately in our case, the prominent cepstral peak is relatively unambiguous in terms of its timing and phase characteristics.

Despite the limitations, the finding of increased inflation volatility in the inflation targeting regime indicated by the positive delta function in the cosine-squared cepstrum cannot be taken lightly. In our opinion, the possible explanation for the increased fluctuations in the CPI inflation lies in the width of the target band of 3 percent to 6 percent. Mishkin (2003) points out that "the use of target bands has a dangerous aspect. Floors and ceilings of bands can take on a life of their own in which there is too great a focus on the breach of the floor or ceiling rather than on how far away actual inflation is from the midpoint of the target range." Too much focus on breaches of the floors or ceilings, in turn, can lead to the so-called instrument instability problem (Bernanke et al. (1999) and Mishkin and Schmidt-Hebbel (2002)). Moreover, it can also lead to suboptimal setting of monetary policy and controllability problems resulting in the inflation target to be missed in the medium-term Mishkin (2003). In this regard, as pointed out by Demertzis and Viegi (2006, 2007, and 2008), a narrower, and possibly also a lower, target band could be of immense help in improving the central bank's credibility and causing inflation expectations to converge to a focal point, and hence, bring down the mean and variance of the inflation rate.