Discussion

Gregor Smith

In this study, Black, Macklem, and Rose thoroughly describe the outcomes when a wide range of reaction functions are set in CPAM, the Canadian Policy Analysis Model developed at the Bank of Canada. The reaction functions for monetary policy influence the short-term interest rate relative to the long-term rate. The analysis begins with the CPAM reaction function, which targets the expected difference from the target of inflation rates seven and eight quarters ahead. The authors then examine generalizations of this rule with different leads and speeds of reaction, rules that use only contemporary indicators, and rules that target the price level.

I draw the reader's attention to four among many noteworthy aspects of the study:

1. The simulation model is ingeniously calibrated so that five-quarter impulse response functions match with those in the data.
2. Monetary policy is separated into a rule (a reaction function for *rsl*) and historical discretion ξ^{rsl} in CPAM.
3. The proposed new rules are assessed using the historical shock properties.
4. The mechanism behind the EF-IP rules, which target price-level stability as well as inflation, is well described in Section 5.

Before they use the model to assess policies, Black, Macklem, and Rose study its properties. In Table 2 they document a wide variety of statistical moments implied by the model with the CPAM reaction function. They note that the model's moments resemble those found in recent historical data. To me, this test is not completely convincing. After all, since the CPAM reaction function differs from the policy rules that were followed historically, one might well expect various moments to be affected by it. In fact, the effect appears mainly in the variance of inflation (which is lower in the model than in the data) but not in other moments.

Nevertheless, this comparison is interesting. In fact, one could go further and use CPAM to produce counterfactual historical sample paths (not just moments) given the historical shocks. That would show how the economy would have reacted, quarter by quarter, to the CPAM reaction function, of course assuming that other aspects of the calibrated model would have been unchanged so that the Lucas critique would not apply.

The next step in the investigation is to study the volatilities of output and inflation under different reaction functions. The paper uses graphs of these two volatilities to describe the results of an enormous number of simulated policy scenarios. One striking finding is that EF-I rules can do better than C-IY rules.

In my role as sceptic, I wonder whether the EF-I rules do so well because they are based on model-consistent forecasts of inflation, which of course are correct in the laboratory of the simulations. How much of the policy rules' success comes from the information available to the central bank in forecasting inflation? And how much comes from the bank's presumed superhuman knowledge of the structure of the economy? A slightly different way of posing this question is to ask: How well would EF-I rules do if the environment were not just as supposed in the rest of the model? Would the C-IY rules perhaps be more robust to misspecification of the rest of the model? These questions do not necessarily lead to a clear path for further research, since there are many possible wrong models. But one way to proceed might be to study a *given* reaction function under a range of possible parameterizations of the *rest* of CPAM. One would hope that a policy rule would do well under a range of models, and not hinge on a single property that might be contentious empirically. This is the opposite of what Black, Macklem, and Rose have done so thoroughly, and it might complement their findings.

The best policy rules in their study target forecasts of inflation, which is exactly what happens in inflation targeting in several countries. But an equally important part of inflation targeting seems to be communicating policy to participants in the money and foreign exchange markets and to the public. In Canada this communication is now done through the monetary conditions index (MCI) and through the behaviour of the overnight interest rate. It would be most interesting to see how the MCI and the overnight rate behave along a simulation, since the Bank of Canada's ability to translate a policy rule into these terms may be an important part of targeting.

From the perspective of macroeconomic researchers, the good news in Black, Macklem, and Rose's work is that dynamic, general-equilibrium macroeconomic models can help in designing a good policy rule. My reaction to this good news is: If the shoe fits, wear it. But in case the shoe does not fit, I next will fit an autoregressive conditional heteroscedasticity (ARCH) model. In other words, before agreeing completely that the simulations in their model have identified some good policy rules, I investigate several of its predictions using some time-series models that feature ARCH.

First, the ARCH model can be used to study the link between the output gap and volatility. One of the elements of the CPAM model is a convex Phillips curve, so that low volatility in the output gap is associated with a high value for output. This characteristic of the model influences the design of the policy rules. As a simple test for evidence of this link, I fit an ARCH-in-mean model to the output gap, with the gap measured just as in the paper. The model is:

$$y_t = \beta_0 + \beta_1 y_{t-1} + \beta_2 \sqrt{h_t} + u_t$$
$$u_t \sim N(0, h_t)$$
$$h_t = \alpha_0 + \alpha_1 u_{t-1}^2 ,$$

where y is the HP-filtered, log real quarterly GDP (D20463), from 1947Q4 to 1996Q4. This statistical model uses no information on inflation. Instead it directly studies the correlation between the level of the output gap and the conditional variance of the gap, h.

For the entire sample period, the estimate of β_2 is negative with a t-statistic of –0.2. For the period since 1980, the sign is positive, with a t-statistic near 1.0. This recent point estimate suggests that output gap volatility is positively associated with the gap itself, which is the opposite of the effect hypothesized in the paper. From this test there is no statistical evidence of a stable link between the level and the volatility of the output gap. If there is indeed no link, then the paper may overemphasize the gains from rapidly stabilizing output gaps. Focussing on stable inflation may be a better strategy.

Second, ARCH modelling can also be used to study the empirical evidence on the alleged trade-off between output volatility and inflation volatility. I fitted univariate ARCH models to the output gap and to CPI (P700000) inflation, in annual data. Here the variance is a squared forecasting error, rather than a squared targeting error. Lags in the conditional means are included so that the forecasting models are sensible, and a first-order ARCH model is fitted to each series.

There are three reasons to be interested in this evidence: First, it allows us to compare historical outcomes in different years.

Second, we can use it to test whether historical policies lie on a downward-sloping line, as if on the policy frontier for a given reaction function (this is not a test of whether the frontier is convex for a given rule). It seems worthwhile reporting some evidence on this "variance" Phillips curve, given the extensive study of the standard curve in levels.

Ultimately, I should like to compare the outcomes found by Black, Macklem, and Rose with the historical record. Past outcomes may provide a measuring stick for comparing the different simulated policy rules.

Figure 1 shows the values of the conditional standard deviation of inflation plotted against the conditional standard deviation of the output gap. The graph appears to be convex, but there is not much evidence of a simple trade-off. Outliers are labelled with the year in which they occurred. For

Figure 1

Volatility Trade-Off, 1951 to 1996

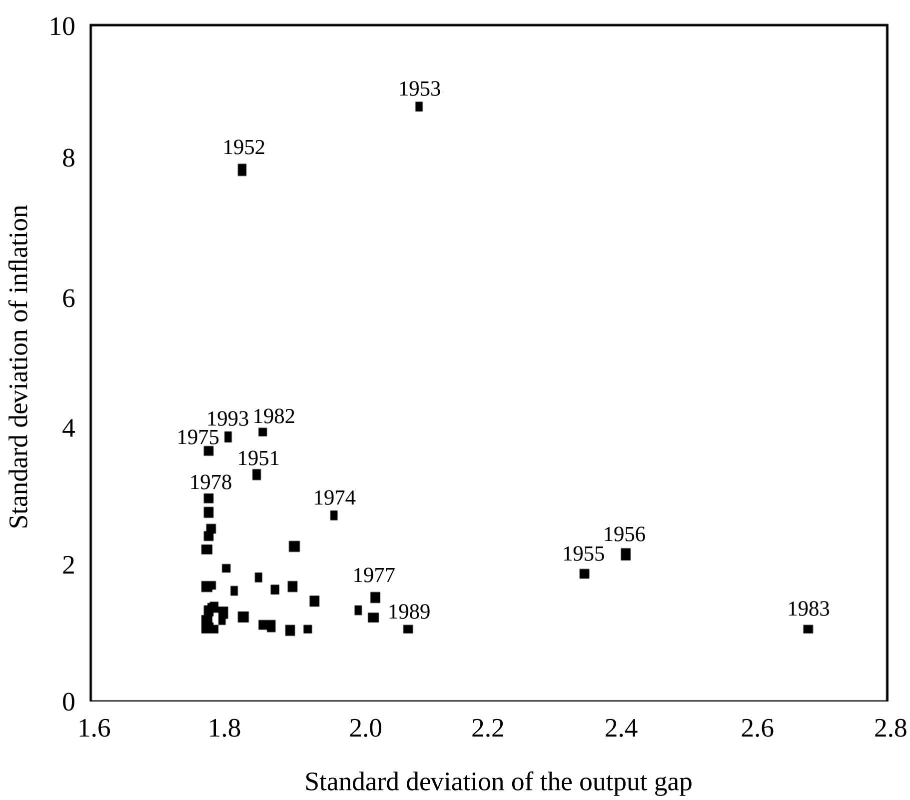

Source: Author's estimates from ARCH models of inflation and the output gap.

example, inflation forecast errors were large during the early 1950s in the disinflation that followed the Korean war.

Figure 2 zooms in on the lower-left corner of the graph. Most recent years are found here, where inflation and output volatility are both low. The exceptions are 1983 and 1993, which both are visible in Figure 1. Again in Figure 2 there is no obvious trade-off, but the shape is convex.

I constructed Figures 1 and 2 using univariate time-series forecasts. A different way to study the volatility trade-off would be to use a panel of forecasts, and estimate the variance using the cross section of forecast errors rather than the ARCH model used here. Exactly this method could be used on the data presented by Johnson (in this volume), which include forecasts

Figure 2

Volatility Trade-Off, 1982 to 1996

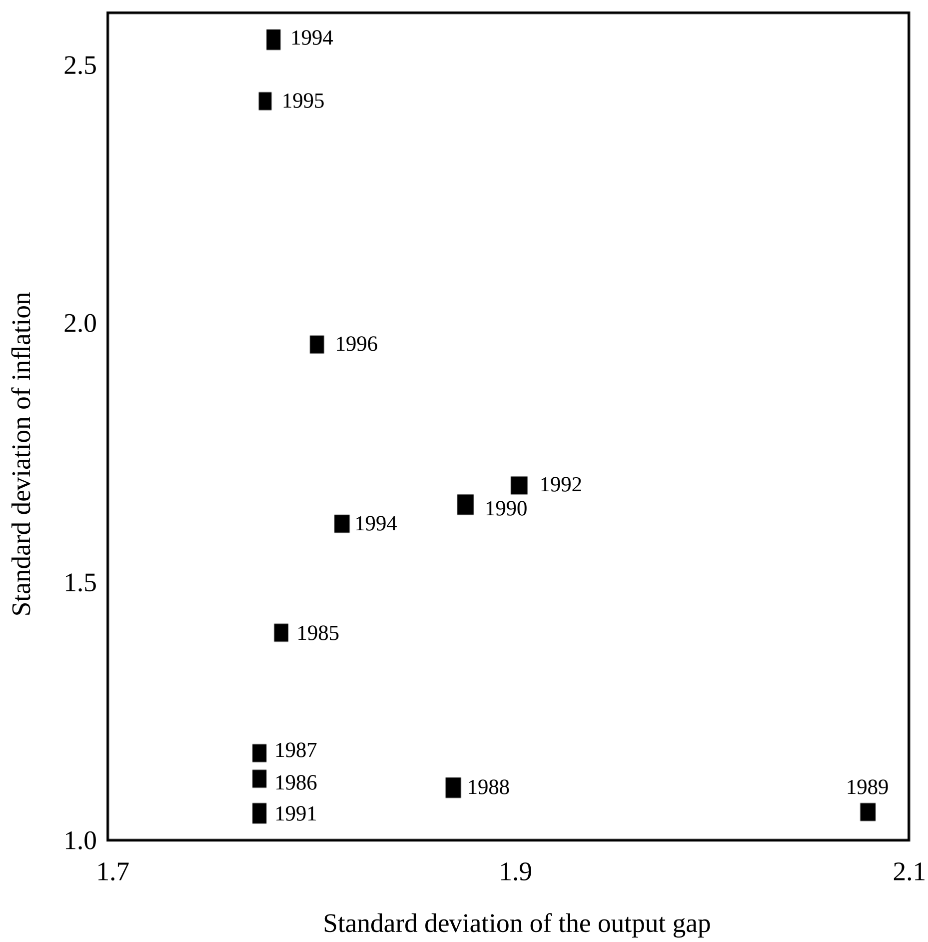

Source: Author's estimates from ARCH models of inflation and the output gap.

of both output and inflation. Countries could be compared using both variances.

My last comment on the study by Black, Macklem, and Rose is in the same spirit of empiricism. Section 5 of their paper studies rules that include a price-stability target, and argues that this target anchors expectations and so helps with stabilization policy. A similar argument has been made by Coulombe (this volume). Their arguments seem speculative to me, for there is considerable historical experience with price-level targeting. An example is found in the gold-standard era discussed by Coulombe. One of the rules of the game under the classical gold standard was that a nation could leave the gold standard during wartime and then return to it later. The United Kingdom did exactly that during the 1914-18 war. In this case the price-level

Figure 3

U.K. Prices and Unemployment, 1920 to 1926

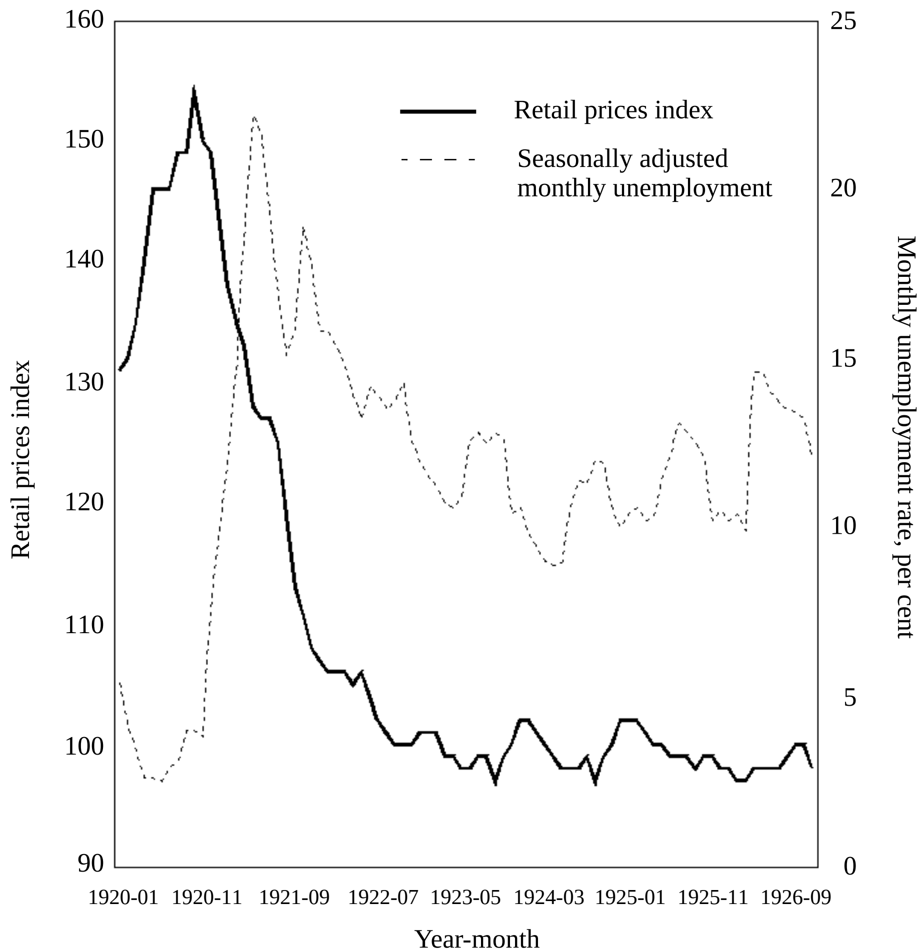

Source: Capie and Collins 1983, Tables 2.14 and 4.5.

target was widely understood, and had been returned to previously, so one would assume that its credibility was high.

Data for this period are available from the statistical study of the United Kingdom between the two world wars by Capie and Collins (1983). The solid line in Figure 3 shows monthly data on the U.K. retail prices index from 1920 to 1926 (from their Table 2.14), seasonally adjusted, with the scale on the left-hand vertical axis. The decline represents the return to the gold standard, which was achieved by 1925. This decline was not an uncertain deflation like those of the early 1930s in some countries, for the target was well understood and had been met during the 19th century.

The dashed line in Figure 3 shows the monthly unemployment rate (from Capie and Collins 1983, Table 4.5), seasonally adjusted, with the scale on the right-hand vertical axis. A quarterly index of industrial production follows a similar time path. The depression of the early 1920s in the United Kingdom was a sharp and disastrous one. In my opinion, we need to study episodes like this one carefully before embarking on a strategy of price-level targeting.

References

Capie, F., and M. Collins. 1983. *The Inter-War British Economy: A Statistical Abstract.* Manchester: Manchester University Press.

Coulombe, S. 1998. "The Intertemporal Nature of Information Conveyed by the Price System." Bank of Canada. This volume.

Johnson, D. 1998. "The Credibility of Monetary Policy: International Evidence Based on Surveys of Expected Inflation." Bank of Canada. This volume.