On Policy Rules for Price Stability

Richard Black, Tiff Macklem, and David Rose*

Introduction

Over the 1980s and 1990s, monetary policy in most industrialized countries has focussed increasingly on achieving rates of inflation at or close to price stability. The result in the 1990s has been a remarkable convergence in these economies towards relatively low rates of inflation. In some countries, such as New Zealand, Canada, the United Kingdom, and Sweden the move towards price stability has been accompanied by the announcement of explicit inflation targets or target ranges.[1] Other countries, such as Germany and Switzerland, have explicit medium-term inflation objectives; still others, such as Japan and the United States, have made price stability, albeit unquantified, a medium-term objective. It is evident that, to varying degrees, monetary policy in these countries is now aimed at keeping inflation low and relatively stable.

In this paper, we address several questions about the conduct of monetary policy in this setting. Specifically, we consider different policy rules for a monetary authority that is committed to price stability, and we examine how these different rules would affect the dynamic properties of the economy. The analysis is conducted using stochastic simulations of a model

1. Finland, Australia, Spain, and Israel have also adopted explicit inflation targets.

* *We are grateful for the valuable technical assistance of Advin Pagtakhan and Hope Pioro. We thank Bob Amano, Paul Fenton, David Longworth, John Nelms, and Simon van Norden for helpful comments.*

of the Canadian macroeconomy, which we call the Canadian Policy Analysis Model (CPAM).

CPAM is a one-domestic-good, small-open-economy description of the Canadian economy, which features forward-looking, though not entirely model-consistent, expectations, an endogenous supply side, behavioural equations for the principal components of demand, and reaction functions for both the monetary and fiscal authorities.[2] It has been calibrated to have the same steady state as the Bank of Canada's Quarterly Projection Model (QPM) as well as very similar dynamic properties for a number of standard deterministic shocks.[3] CPAM is smaller, however, and has been configured to simulate much faster than QPM so that stochastic simulations on the scale required by this project are feasible. The stochastic simulations are implemented in the model using six types of disturbances—shocks to (1) potential output, (2) world commodity prices, (3) the sum of consumption and investment, (4) the price level, (5) the real exchange rate, and (6) the slope of the yield curve. The joint stochastic structure of these shocks is determined using an estimated vector autoregression (VAR) model.

The policy rules we consider all share the feature that a short-term interest rate is the instrument of monetary policy.[4] In general terms, rules are distinguished from discretion by the requirement that following a rule precludes period-by-period reoptimization on the part of the monetary authority. Rules therefore involve precommitment, but this can involve a precommitment to contingent behaviour. The rules we consider are relatively simple in the sense that the contingent behaviour they embody is a linear function of one or two macroeconomic variables. In particular, rather than making assumptions about the preferences of policymakers and deriving optimal rules, we consider relatively simple rules, and then map out the range of possible outcomes under different parameterizations of these rules.[5] Policymakers can then apply their own preferences to determine the most desirable outcome within the set of possibilities.

The restriction to simple rules reflects two considerations. First, optimal rules depend on the particular objective function, and we know too little to pick even the form of such a function, let alone its specific

2. CPAM is described in the next section. See Black and Rose (1997) for details.

3. QPM has been documented in a series of Bank of Canada technical reports. See, in particular, Black, Laxton, Rose, and Tetlow (1994) and Coletti et al. (1996).

4. This reflects the reality of modern monetary policy. The preferred instrument among central banks in industrialized countries is a short-term interest rate.

5. For an alternative methodology, where the optimal control problem for a linear model is solved for a continuum of possible calibrations of quadratic preferences for policymakers, see Tetlow and von zur Muehlen (1996).

parameters. Second, once one moves beyond linear models, optimal rules depend, in general, on the details of the structure and parameters of the underlying economic model and on the characteristics of the shocks. Indeed, optimal control results cannot, in general, be achieved with any explicit (that is, closed-form) rule. Given that an important goal of our research is to find a characterization of contingent behaviour that is simple enough that it could be effectively communicated to private agents as part of the commitment to a policy of price stability, we view optimal rules as impractical. In this regard, our research contributes to recent attempts to find simple rules that provide reasonable results over a wide range of shocks. See, for example, the studies in Bryant, Hooper, and Mann (1993).

We begin by examining the stochastic behaviour of macro variables of interest in the model when the monetary authority uses the base-case reaction function in CPAM.[6] This reaction function calls for the monetary authority to act on the basis of the *expected* deviation of inflation from its targeted rate seven to eight quarters into the future. The targeted inflation rate is taken to be 2 per cent a year—the midpoint of the announced inflation control range in Canada of 1 to 3 per cent. We begin with questions about the basic properties of the stochastic economy, such as:

- What are the implications of this forward-looking, inflation-targeting reaction function for the distributions and dynamic properties of key macro variables, such as inflation, output, interest rates, and the exchange rate?
- What proportion of the time is inflation expected to be outside the bands of 1 to 3 per cent, and what is the average duration of the departures from these bands?

The answers to these questions serve as a benchmark of comparison against which to measure the performance of alternative monetary policy rules, including alternative calibrations of the base-case reaction function.

We consider three broad classes of policy rules: (1) reaction functions, like those in CPAM and QPM, that involve adjusting interest rates to the expected future deviations of inflation from its target level; (2) reaction functions in which interest rates respond both to deviations of inflation from its target and to deviations of output from potential, but using only contemporaneous information; and (3) reaction functions that call for the monetary authority to return the price *level* to a target path.

We refer to the first class of rules as explicitly forward-looking (EF), in that the response is to expected future outcomes. The particular subclass we study here is limited to rules where only the deviation of inflation (I)

6. This reaction function is patterned on the one used in the Bank of Canada's QPM.

from the target is given weight in the instrument setting; we therefore use the label EF-I for this subclass.

The second class of reaction functions has some similarity to what has come to be called the Taylor rule, following, for example, Taylor (1993, 1994).[7] Our implementation is not quite the same as Taylor's, however, and we label this class as C-IY, for rules with response to contemporaneous (C) values of inflation (I) and output (Y).

The third class of rules we consider is another subclass of EF rules. In EF-I rules, only the rate of inflation matters. When there is a shock that moves inflation away from the target, the authority works to return inflation to that target level, but no consideration is given to eliminating the dynamic drift in the price level away from the constant-inflation path during the disequilibrium. Thus, even though the monetary authority may succeed in anchoring inflation to the target level, the deviation of the price level from a constant-inflation path will have a unit root (the log price level will not be trend stationary). The third class of rules we consider adds the explicit imposition of trend stationarity for the price level.[8] We call this class EF-IP, for inflation plus a condition on price (P) drift. Note that we are not talking about imposing a stationary price level—literal price stability—but rather the weaker condition that there is no cumulative drift in the price level away from a control path of constant inflation.

Taylor (1979, 1994) argues that, while policymakers do not face a permanent trade-off between the level of inflation and the level of the output (or unemployment) gap, they do face a long-run trade-off between the variability of inflation and the variability of output. We examine this trade-off by generating efficiency frontiers for reaction functions in the EF-I and C-IY classes. We also conduct a more limited examination of reaction functions in the EF-IP class, and compare their performance against outcomes along the jointly efficient frontier for reaction functions in the EF-I and C-IY classes.

---

7. Taylor is not the only researcher to study rules with this form. Henderson and McKibbin (1993) study various rules, including one with a similar structure. Taylor's particular choice of parameters for his rule, which he selects as a representation of historical U.S. monetary policy, imposes the same coefficient on the inflation and output gap terms. This makes it, in essence, a nominal income rule, of which there are many examples. See Bryant, Hooper, and Mann (1993) for a thorough discussion and evaluation of similar policy rules, and Levin (1996) for a recent evaluation of the Taylor rule and his version of the Henderson-McKibbin rule using the Federal Reserve Board's Multi-Country Model.

8. We do not mean to imply that small-sample statistical tests will always indicate trend stationarity. Nevertheless, the policy response imposes a condition that the price level is expected to return to a given constant-inflation path.

There are also issues that go beyond pure variability. In the linear framework of Taylor's simple model, certainty equivalence properties emerge—in a stochastic environment the means of the distributions will generally be invariant to the choice of a rule, as long as the policy rule is stable to begin with.

The model we use is not linear. In particular, motivated by recent empirical evidence, we have configured it to feature an asymmetry in inflation dynamics, whereby excess demand creates more upward pressure on inflation than an equivalent degree of excess supply creates downward pressure. In this case, certainty equivalence does not hold, and we must also consider results for first moments. In our analysis it is of particular interest whether a given rule achieves the inflation target on average, and whether there is any associated shift in the level of output that is attainable on average over time. Finally, we are dealing with an open economy, and we must therefore also consider the volatility of the exchange rate and interest rates.

We address several questions, including the following:

- Along the efficient policy frontier for reaction functions in the EF-I class, what is the slope of the trade-off between inflation and output variability, and what are the implications for the variability of interest rates and the exchange rate of moving along this frontier? Are the means of the distributions of the main macro variables affected? What calibrations of this class, in terms of horizon for the control lead and for the parameters of response, produce efficient outcomes?
- How does the efficient policy frontier for the C-IY class of reaction functions compare with that of the EF-I class? What weights on the inflation and output terms produce efficient outcomes? Are there other consequences for the levels or variability of the main macro variables?
- Do reaction functions of the EF-IP class, which add an element of price-level control, unambiguously underperform reaction functions based on pure inflation targeting? Or can a case be made for an element of price-level targeting, based on its implications for the expected levels or variability of the main macro variables?

The paper is organized as follows. Section 1 provides a brief overview of the model, highlighting certain key features that are important for the results. Section 2 describes the VAR methodology we use to estimate the stochastic structure of the shocks, with details in Appendix 1. Section 3 presents the results using the base-case CPAM reaction function, while Section 4 examines the prospects for improved inflation control, and the possible consequences in terms of other properties, for the EF-I and C-IY classes of reaction functions. Section 5 adds the extension to price-level targeting provided by the EF-IP class. Finally, we offer some concluding remarks.

1 The Canadian Policy Analysis Model, CPAM

This section provides an overview of CPAM, the model we use for our stochastic simulation analysis. We do not give a detailed description of CPAM in this paper, but focus instead on its broad features, as well as certain key relationships that are important for the analysis here. A complete description of the model and its properties in response to various standard deterministic shocks can be found in Black and Rose (1997).

CPAM is designed to provide a reasonably complete representation of the Canadian macroeconomy.[9] It has quarterly frequency and is made up of about 140 equations, of which perhaps 30 describe the essential agent behaviour. The model has an explicit steady state and is dynamically stable over a wide range of disturbances. It accounts fully for three stocks—production capital, government bonds, and net foreign assets—and the steady state describes an equilibrium of a small open economy with stable ratios of each of these stocks to output.

CPAM has many features similar to the Bank of Canada's QPM, and it has been calibrated to reflect current Bank of Canada staff judgments regarding exogenous variables, the numerical steady-state solution, and many features of dynamic properties in deterministic simulations. However, CPAM is smaller and is configured to simulate much faster than QPM, which is essential for the stochastic analysis we provide in this paper.

CPAM represents a small open economy that produces a single domestic good. There are four groups of domestic agents in the model: firms, consumers, and the fiscal and monetary authorities. Profit-maximizing firms combine labour and capital in a Cobb-Douglas technology to produce the single good. Trend population growth and trend productivity growth are exogenous.

Consumers come in two types. There are forward-looking consumers who make decisions with a view to picking the best path for current and future consumption, and "rule-of-thumb" consumers who spend all their income in each period. Thus, all assets are held by the forward-looking agents. The behaviour of the forward-looking consumers is characterized

9. CPAM was prepared for the Bank of Canada by David Rose at QED SOLUTIONS, with input from Richard Black on the tuning to QPM properties. CPAM exploits several features from work done by QED SOLUTIONS for the Reserve Bank of New Zealand, in the development of a new model of the New Zealand economy. This work will be documented in a forthcoming Reserve Bank publication on its new Forecasting and Policy System.

using the Blanchard (1985)-Weil (1989) model of overlapping generations, but in a discrete-time format, as in Frenkel and Razin (1992). CPAM uses a version of this approach with some adaptations to speed up the adjustment to asset equilibrium.

There is a single consolidated federal-provincial-municipal fiscal authority. Fiscal policy is characterized as the choice of target ratios for debt to output, for spending on goods and services to output, and for transfers to the household sector to output. There is an indirect tax on consumption and a personal direct tax. The indirect tax rate is treated as an exogenous fiscal choice. The fiscal authority's intertemporal budget constraint is respected, and the fiscal targets met, through a reaction function that sets the personal direct tax rate endogenously.

Monetary policy is characterized as a choice of a target inflation rate for consumer prices. This target is achieved through a monetary reaction function, where the instrument is the short-term (90-day) interest rate, but where this instrument influences the real economy through its effect on the slope of the term structure. We describe this reaction function in more detail below.

The domestic absorption deflator at factor cost serves as the numeraire for the model's accounts and is the core measure of aggregate domestic prices in the model. All other price indexes are built from this core price and external trade prices. Thus, for example, the consumption good price in the model is determined as a weighted combination of the import price and the domestic price for consumption goods, all appropriately adjusted for indirect taxes, using the model's accounting identity structure. In the model's small-open-economy paradigm, the foreign price is exogenous, and the law of one price holds on the margin. The Phillips curve in CPAM relates core inflation to domestic demand conditions. But much of consumer price dynamics comes from the direct effect of import prices on consumer prices. The exchange rate therefore plays a central role in the monetary transmission mechanism.

The steady-state *real* exchange rate in CPAM is determined endogenously to support the overall asset equilibrium. Given a supply of labour and relative factor prices, an equilibrium stock of capital is determined from the conditions for profit maximization. The stock of government bonds comes from the fiscal policy choice of a debt-to-output ratio. The household intertemporal optimization then provides the residual asset level, the stock of net foreign assets (which, for Canada, is a negative value—that is, Canada is a net foreign debtor and the model is calibrated to reflect this fact). To support the net foreign asset position, and the consequent debt service flow, a particular level of surplus is required for the trade balance. Essentially, the steady-state real exchange rate is determined to support this equilibrium trade balance.

In the short run, however, the real exchange rate responds to interest rates, and to monetary policy actions in particular. In the steady state, domestic real interest rates are tied to exogenous foreign values, but in the short run they are determined by a version of the standard uncovered interest-parity condition.

We now turn to selected aspects of the model that require a more elaborate exposition for the purposes of this paper.

1.1 The reaction function

The general EF-I reaction function in CPAM is written as follows:

$$rsl = rsl^* + 1.6 \bullet (\pi c_4(7) - \pi^*) + 1.8 \bullet (\pi c_4(8) - \pi^*) + \xi^{rsl}, \qquad (1)$$

where rsl is the slope of the term structure of interest rates (the short rate relative to the long rate), rsl^* is the steady-state value of rsl (which is constant in our analysis), πc_4 is the four-quarter rate of change of consumer prices, and $x(f)$ indicates an f-quarter lead of the variable x, which is resolved as a model-consistent solution in simulation.[10] In our experiments, the monetary authority has a fixed target for the inflation rate, given by π^*.[11] This is set at 2 per cent a year, the midpoint of the current target range for inflation in Canada. Finally, ξ^{rsl} is a shock term that captures other influences on the yield spread, such as the effects of changes in the slope of the U.S. term structure or changes to the risk premiums embedded in long rates.

To hit the inflation target, the monetary authority acts to move the intermediate target variable, rsl, raising (lowering) relative short rates when inflation is expected to be above (below) the target rate seven or eight quarters ahead. The desired outcome for rsl is achieved by setting the true instrument variable—the short-term nominal interest rate—at the required level. The model contains a description of the dynamics of the term structure of interest rates, in both real and nominal terms. Thus, the short nominal rate instrument is adjusted taking into account any movements in long rates in order to achieve a particular desired setting for rsl, and ultimately for the expected path of inflation.

10. In CPAM, the long rate is formally a 5-year rate, but the model is calibrated to roughly mimic the properties of QPM, where the long rate is for a 10-year instrument. This limitation of the horizon is one of the concessions we made in CPAM to speed simulation times.

11. We also assume that the economy has already adjusted to this target. There are no “transition” issues; rather, we look at conditions that might emerge from various approaches to sustaining the target rate.

The use of *rsl* in the reaction function instead of simply the level of the real short rate follows QPM and reflects several considerations. For our purposes, the main one is the following. Historical movements in both long and short rates reflect changes in the equilibrium real interest rate as well as monetary policy actions and inflation expectations. Moreover, short rates reflect changes in the stance of monetary policy more strongly than long rates; thus to a large extent the spread serves to isolate the monetary component of real interest rate changes. This is helpful for the identification of monetary and non-monetary shocks in the historical data (which is important for our stochastic analysis as described below).[12]

1.2 The Phillips curve

The CPAM Phillips curve proximately determines the rate of inflation of the model's numeraire price, the domestic absorption deflator at factor cost. Since the central bank is presumed to formulate its actions on the basis of the anticipated course of consumer prices, and since a significant proportion of consumption goods are imported, the dynamics of domestic prices are just part of the overall inflation story. Nevertheless, the Phillips curve is central to the monetary transmission mechanism in the model.[13] The essence of the model equation is as follows:

$$\pi = 0.25\bullet(0.6\bullet\bar{\pi}^{e} + 0.4\bullet\bar{\pi}c^{e}) + 0.43\bullet\pi(1) + 0.32\sum_{i=1}^{4}\alpha_i\pi(-i)$$
$$+ 0.02\bullet ygap + 0.04\bullet ygap(-1)$$
$$+ 0.09\bullet ygap^{+} + 0.03\bullet ygap^{+}(-1) + \xi^{\pi}\,. \qquad (2)$$

The CPAM Phillips curve has many standard features. It imposes the long-run natural-rate hypothesis (that is, there is no permanent trade-off between output or employment and the rate of inflation). In the short run, however, there is a dynamic link between excess demand, *ygap*, and

12. For a discussion of the yield curve as an indicator of monetary stance, see Côté and Macklem (1996) and Macklem, Paquet, and Phaneuf (1996).

13. The actual equation has one other important term that replaces other structure in the larger model from which CPAM was drawn. This term operates only when there is a shock to the real steady state, which for us means the pure supply (productivity) shock. It increases the downward pressure on prices from a positive supply shock, reducing the share of real wage adjustment that comes through nominal wages.

inflation, π.[14] Recent empirical evidence suggests that this linkage is asymmetric, as in the original Phillips (1958) curve. This asymmetry is such that the positive effect of excess demand on inflation is stronger than the negative effect of an equivalent degree of excess supply. For Canada, results documented in Laxton, Rose, and Tetlow (1993a) led to this feature being included in QPM. Turner (1995) and Debelle and Laxton (1996) have also found significant asymmetry of this sort for the Canadian data.[15] The form of this function in CPAM is the same as in QPM—a piecewise linear version with a steeper slope when excess demand is positive, as provided by the $ygap^{+}$ terms in equation (2).

The first part of equation (2) captures intrinsic and expectational dynamics. The structure is based on a contracting paradigm, with periodic bargaining, as in Fuhrer and Moore's (1995a, 1995b) "real wage" version of Taylor's (1980) model. We assume annual bargaining, so the terms $\bar{\pi}^{e}$ and $\bar{\pi}c^{e}$ are averages of expectations formed in the current and previous three recent quarters when contracts still extant were signed. Thus, for example,

$$\bar{\pi}^{e} = 0.25 \sum_{i=0}^{3} \pi^{e}(-i) , \tag{3}$$

where $\pi^{e}(i)$ is the expectation that was formed i quarters in the past. We describe how expectations are formed in the next subsection. The equation for $\bar{\pi}c^{e}$ has the same form. The presence of both expectations reflects the view that both consumers and firms have some influence in the wage-bargaining process, with firms caring more about the real wage in producer prices, and consumers caring more about the real wage in consumer prices.

---

14. In simulation, the level of potential output in the model comes from evaluating the production function at full employment and trend factor productivity, but with the existing stock of capital. This treats capital as a quasi-fixed factor and employment as completely variable in the short run, as is traditional in such models.

15. Dupasquier and Ricketts (this volume) also find evidence of this type of asymmetry in the Canadian Phillips curve, while noting that this finding is sensitive to how the output gap is measured. Evidence of this type of asymmetry in price adjustment has also been reported for other countries. Turner reports significant asymmetry for the United States and for Japan. Laxton, Meredith, and Rose (1995) find strongly significant asymmetry for a pooled sample of G-7 countries. Clark, Laxton, and Rose (1996); Debelle and Laxton (1997); Clark and Laxton (1997); and Dupasquier and Ricketts (this volume) find significant asymmetry for the United States. Fisher, Mahadeva, and Whitley (1996) find significant asymmetry for the U.K. data, as do Debelle and Laxton. Isard and Laxton (1996) find significant asymmetry for a pooled sample with France, Italy, and the United Kingdom. Bean (1996) reports evidence for modest asymmetry for a panel of OECD countries.

The $\bar{\pi}^e$ and the $\bar{\pi}c^e$ terms are thought of as reflecting cost pressures, but that is not the only source of inertia. The term $\pi(1)$ represents a one-quarter-ahead, model-consistent forecast, and there are also lags of the actual inflation rate (with $\sum \alpha_i = 1$) reflecting quarterly price adjustment by firms conditional on the underlying cost trend.

1.3 Expectations of inflation

Inflation expectations are modelled using a variant of the Buiter-Miller (1985) mixed model with both forward- and backward-looking components. A typical equation is:

$$\pi^e = 0.35 \bullet (0.7 \bullet \pi(-1) + 0.3 \bullet \pi(-2)) + 0.55 \bullet \pi_4(4) + 0.1 \bullet \pi^{*e}. \tag{4}$$

The backward-looking component is a weighted average of the previous two observations of inflation. The forward-looking component is a model-consistent forecast. The term $\pi_4(4)$ stands for a four-quarter rate of change; it enters with lead 4, which makes this term the forecast rate of inflation over the next four quarters, the assumed bargaining horizon.

We also add an explicit weight on the *perceived* target inflation rate.[16] This variable is determined as follows:

$$\pi^{*e} = 1.35 \bullet \pi^{*e}(-1) - 0.425 \bullet \pi^{*e}(-2) + 0.0375 \bullet (\pi_4(16) + \pi_4(20)). \tag{5}$$

The perceived target rate of inflation evolves as a second-order transfer function, with the underlying process driven by a weighted average of the model-consistent forecast for the four-quarter inflation rate four and five years ahead. This term has a relatively small weight, 0.1, in the expectations equation (4). It is designed to represent the effects of credibility. If the monetary authority is expected to keep inflation within a reasonable range of the target level in the medium term, then the expected target will remain very close to the announced target. This will provide something of an anchor to expectations, damping their response to short-term cyclical effects. By contrast, if the monetary authority's reaction function is not expected to keep inflation close to the target, this term will pull expected inflation away from the announced target, even if the recent history is good and the short-run prospects are for inflation to remain close to that target level. This

16. This term was introduced initially in anticipation that part of this research would involve such issues as analysis of credibility and learning. In the end, we did not get to those questions.

captures the idea that it can take more than a few good outcomes for the monetary authority to gain credibility, and that credibility is fragile and, once lost, is hard to regain.

2 The Stochastic Specification

Stochastic simulations require a model and a distribution from which to draw the random shocks. In the case of estimated models, the distribution of shocks is usually based on the distribution of the estimated residuals attached to the various equations in the model. For calibrated models, there is no equally obvious approach. Most previous studies using calibrated models to evaluate alternative monetary policy rules have based the distribution of shocks on the properties of the residuals from various estimated reduced-form equations or VARs under the assumption that the shocks are independently distributed. We use a more ambitious approach.

Our method can be described briefly as follows. We begin by estimating an *N*-variable VAR in which the variables are determined by the *N* types of shocks to be considered. The estimated VAR is used to produce impulse responses to innovations in each of these *N* variables. We then use these impulse responses to define the shock terms for the structural model. In the general case, a shock term in the structural model is a combination shock that includes an innovation to each of the variables in the VAR. This combination shock is defined so that the dynamic response in the structural model matches the VAR impulse response over some horizon (which we take to be five quarters).

We consider six types of disturbances, shocks to (1) potential output, (2) world commodity prices, (3) the sum of consumption plus investment (which serves as our demand shock), (4) the price level, (5) the real exchange rate, and (6) the slope of the yield curve. The disturbance terms for the structural model are defined so as to reproduce the impulse responses of a VAR estimated for these six variables. For example, consider the shock to consumption plus investment $(c + i)$. The VAR produces the estimated impulse responses to a one-standard-deviation innovation in $c + i$ for each of the six variables. The $c + i$ shock for the structural model is chosen to replicate the impulse responses of the VAR over the first five quarters following the shock. To do this, a 6 x 5 matrix of coefficients is chosen. The first column applies to the vector of contemporaneous disturbances, one for each of the six variables in the VAR, and the remaining 6 x 4 matrix weights past values of these disturbances. The technical details are provided in Appendix 1.

This VAR approach has two attractive features. First, it allows the statistical properties of the shocks to be considered in a single unified framework. Second, it allows the autocorrelation and cross-correlation structure of the shocks to be determined by their historical correlations so there is no need to make arbitrary assumptions, such as independence. We view this as an advantage, since we do not interpret shocks to endogenous variables (such as $c + i$) as deep structural shocks, but rather as aggregate summary shocks that reflect combinations of many different types of disturbances at the micro level. Since the deep structural shocks (which we cannot uncover) will, in general, affect several variables at once, there is no reason to expect the aggregate summary shocks to be independent. To the extent that the VAR approach does a better job of capturing the correlation structure of the shocks facing the economy, we can have more confidence in the quantitative analysis.

We now turn to a more detailed discussion of the estimated VAR.

2.1 The VAR

The variables in the VAR are potential output (y^*), real commodity prices (*pcom*), the domestic price level (py), the real exchange rate (z), the slope of the yield curve (rsl), and the sum of consumption and investment $(c + i)$.[17] Potential output and real commodity prices are treated as exogenous variables (consistent with CPAM), and shocks to these variables have obvious interpretations. We interpret the shock to the price level as a supply shock that takes the form of a shift in the expectations-augmented Phillips curve. The real exchange rate shock is interpreted as a shock to investor confidence in the value of the Canadian dollar. We interpret the shock to $c + i$ as a domestic demand shock, and view shocks to the yield curve as capturing monetary disturbances. Together these six shocks are designed to capture the principal macro disturbances facing the Canadian economy, though there are certainly other types of shocks that have played a role.[18]

The variables in the VAR are defined as follows. Potential output is measured as the Hodrick-Prescott filter of actual output (using the usual

---

17. We aggregate consumption and investment rather than entering them separately in the VAR so as to preserve degrees of freedom.

18. Fiscal shocks are an obvious omission from our VAR. We do not consider a stochastic element to fiscal policy because we think that the historical distribution of fiscal shocks is unlikely to provide a good description of fiscal shocks over the next decade. A second omission is the absence of a shock to foreign demand for Canadian goods. This reflects both the small open economy structure of CPAM and our inability to identify a sensible export demand shock in the estimated VAR. To the extent that changes in the level of world (particularly U.S.) economic activity affect other variables in the VAR, such as commodity prices and the yield curve, the effects of export demand shocks are likely captured in part in the VAR in the statistical properties of these other variables.

setting for the smoothness parameter of 1600).[19] Real commodity prices are defined as the Bank of Canada commodity price index, defined in U.S. dollars, divided by the U.S. gross domestic product deflator. The domestic price level is measured as the Canadian GDP deflator. To be consistent with the structural model, the real exchange rate is defined as the nominal rate with the G-6 countries times the aggregate price level for the G-6, all divided by the domestic absorption deflator at factor cost. The yield spread is the 90-day corporate paper rate less the yield on 10-year Government of Canada bonds. Consumption in the model is broadly defined to include expenditures on consumer goods and services, housing, and inventories, while investment includes spending both on structures and on machinery and equipment. The sum of consumption plus investment therefore includes all these components. The yield spread enters the VAR system in levels and the other five variables are in log levels. Since potential output and real commodity prices are viewed as exogenous, the equations for these variables include only lags of the dependent variable, one lag in the case of potential output and four lags for commodity prices. In general, the endogenous variables in the system include four lags of all endogenous variables, as well as one lag of potential output and four lags of real commodity prices. The VAR is estimated over the sample 1961Q1 to 1992Q4. The sample ends in 1992Q4 to avoid the end-of-sample problems associated with using the HP filter to define potential output.

The moving-average representation of the VAR is identified by imposing a Wold causal ordering on the variables in the VAR; we use the ordering $\{y^*, pcom, c+i, py, z, rsl\}$, which is motivated by the structure of the model. In particular, potential output and real commodity prices are both exogenous in the model, while $c+i$ and prices both adjust with a lag to changes in monetary conditions, and the monetary authority is assumed to look at the contemporaneous state of the economy before deciding on the monetary innovation. We also impose two over-identifying restrictions on the VAR. First, we restrict that there be no contemporaneous cross-correlation between potential output and real commodity prices, since these are exogenous variables in the model. Second, we impose the restriction that the lags of the real commodity prices do not appear in the equation for the GDP price deflator. This restriction reflects our judgment that the large effects of the oil-price shocks in 1973 and 1979 on inflation have more to do

19. We recognize that the HP filter is a crude measure of potential output (Guay and St-Amant 1996), but its simplicity and transparency are attractive. We also experimented with the Bank's measure of potential output, which adds structural information to an HP filter, based on QPM. The results were generally similar using this alternative, but in a few cases the responses using the simple HP-filter measure were more conformable with CPAM.

with the fact that the Bank of Canada accommodated these oil-price shocks than they do with the strong structural link between relative prices and the aggregate price level. As a first approximation, we therefore turn off any direct effect of relative prices on the aggregate price level. The impact of this restriction is to improve the conformability between the effects of a commodity price shock in the VAR and the model.

Figure 1 reports the impulse responses of the variables in the VAR to each of the six shocks. The shocks are of magnitude one standard deviation, and the variables being shocked are along the diagonal of the matrix of the graphs. The graphs should be read vertically—each column presents the impulse response to one of the shocks. For example, the first column of results shows the responses to an impulse shock to potential output, the second column shows responses to a commodity price shock, and so on.

Overall, the impulse responses for all six shocks provide a surprisingly coherent macro representation of the effects of the main aggregate disturbances facing the Canadian economy, and one that corresponds well with the economic theory underlying CPAM.[20] The shock to $c + i$ is illustrative. Since potential output and commodity prices are exogenous, these variables remain unchanged. Consistent with the interpretation of the $c + i$ innovation as a shock to domestic demand, the rise in $c + i$ causes the price level to begin to move upwards about one year after the innovation. The monetary authority tightens pre-emptively to contain price pressures. This causes the yield spread to rise and the exchange rate to appreciate. This tightening in monetary conditions dampens demand to the extent that demand cycles below its steady state before returning to an unchanged control level. The negative cycle in demand together with the exchange rate appreciation puts downward pressure on prices, unwinding the initial increase.

The decomposition of variance for the endogenous variables in the VAR, shown in Table 1, provides an alternative way to describe the VAR, which also fits well with the structure of CPAM. In particular, consistent with the view of the monetary transmission mechanism in CPAM, there is no direct effect of monetary shocks (as captured by innovations in rsl) on prices. Rather, monetary or rsl shocks affect demand (as captured by $c + i$) both directly and indirectly through the real exchange rate, and demand in turn influences prices. The variance decomposition also provides some information on the relative importance of the different sources of disturbances. Note that the exogenous shocks—potential output and commodity prices—explain about half the variance after six quarters in

20. See Black and Rose (1997) for a discussion of the dynamic effects of these six shocks in deterministic simulations with CPAM.

Figure 1
VAR Impulse Responses

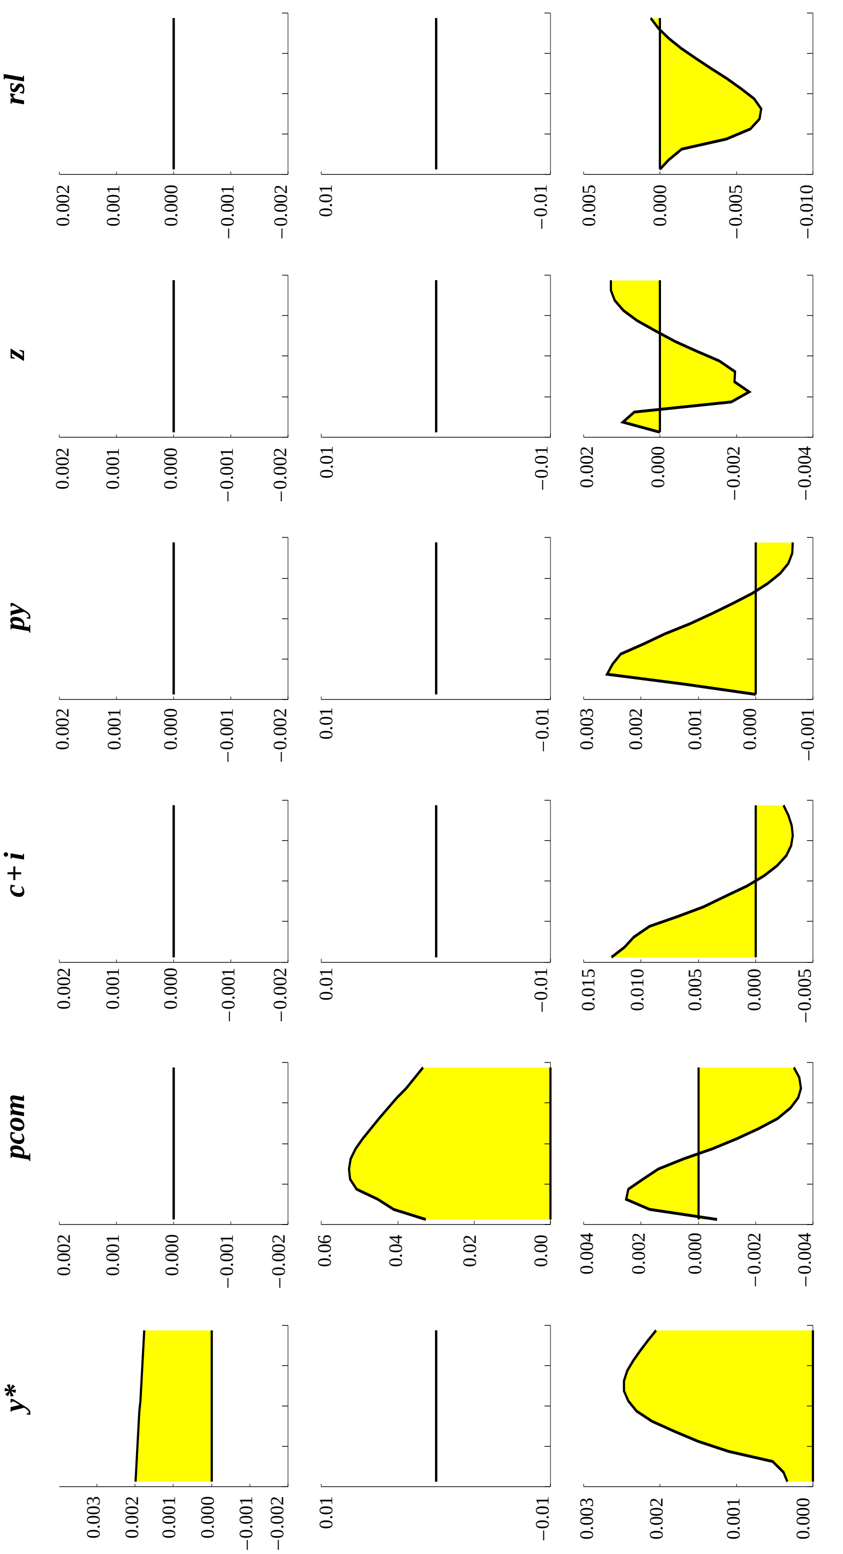

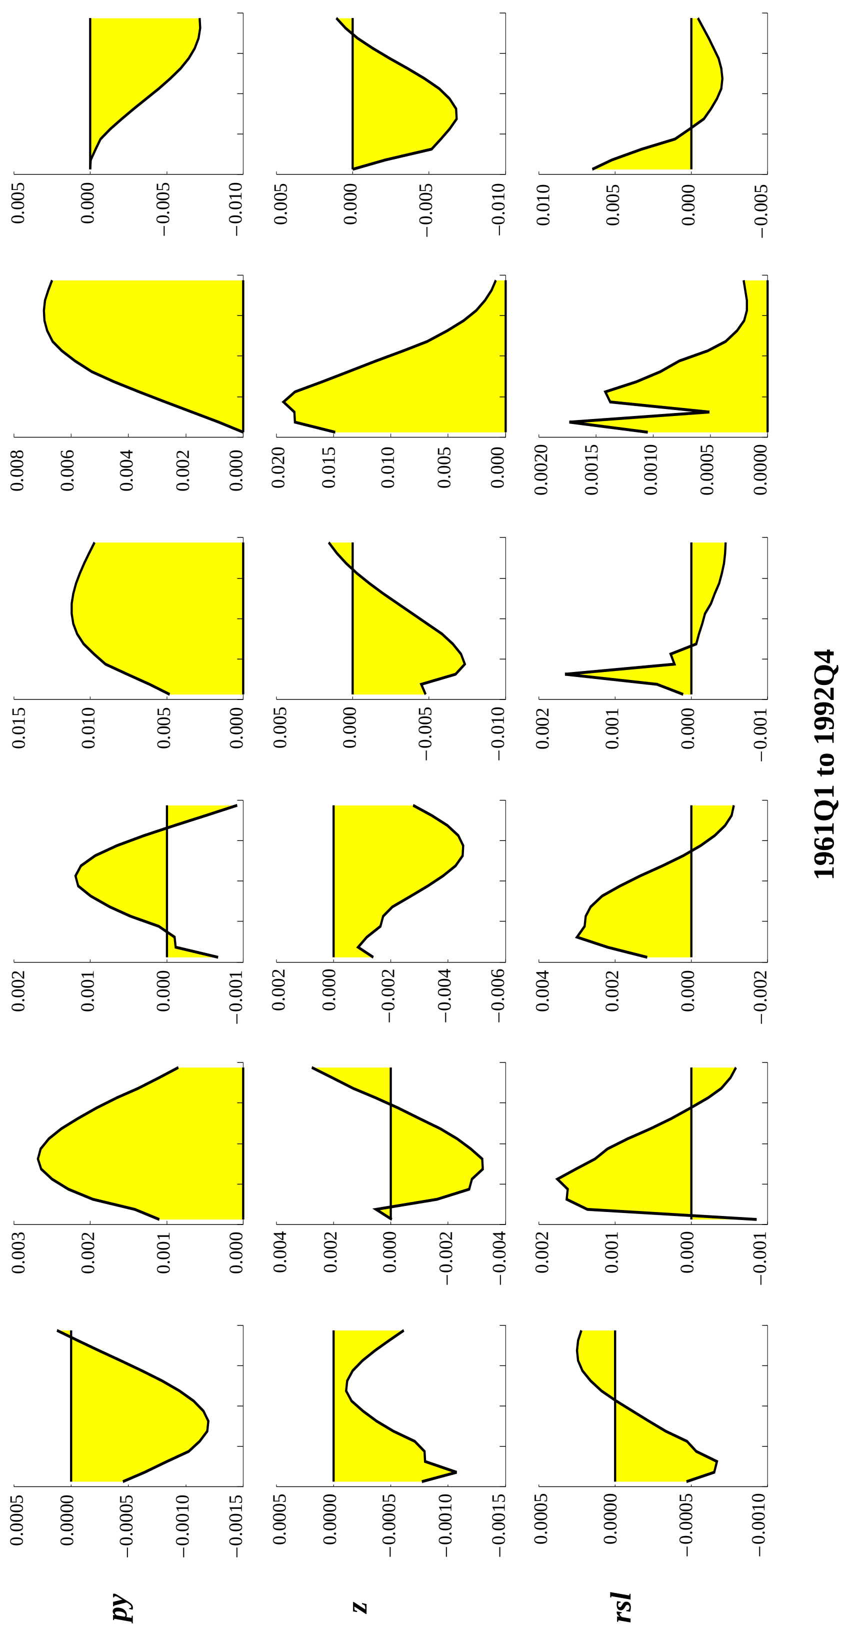
py
z
rsl
1961Q1 to 1992Q4

Table 1

Variance Decomposition for VAR

	Contribution, after six quarters, of innovations in					
Variables	***y****	***pcom***	***c + i***	***py***	***z***	***rsl***
c + i	29.1	22.7	20.9	5.3	11.9	10.2
py	25.5	25.9	26.0	22.4	0.0	0.2
z	16.0	17.6	18.4	15.6	13.3	19.2
rsl	16.5	24.1	21.8	15.8	11.3	10.6

$c + i$ and prices, while for the real exchange rate and the yield spread, the variance decomposition is more evenly spread across all six types of disturbances.

3 The Stochastic Behaviour of the Economy with the CPAM Rule

We begin our analysis of policy rules by examining the stochastic behaviour of the economy under the base-case CPAM policy rule (hereafter simply the CPAM rule). This reaction function is of particular interest because it has the same general form as the rule that is currently used by Bank staff to provide policy advice, and it has been calibrated (along with the rest of CPAM) so that the responses to a standard set of deterministic shocks in CPAM are reasonably similar to those in QPM. A thorough analysis of the CPAM rule also allows us to consider the stochastic properties of the model in some detail before beginning our study of alternative reaction functions.

To examine prospects for inflation control using the CPAM rule, we perform stochastic simulations with CPAM using the shock structure derived from the VAR. The model is run over 136 quarters, and we drop the first 8 quarters; this leaves 128 quarters, which corresponds to the length of the historical sample over which the VAR is estimated. We perform 10 replications of history, so the total number of simulations is 10 x 136 = 1,360.[21] The distributions for variables of interest are built up by averaging across time and across replications.

21. A forward solution must be computed at each time point, conditional on the information known to that point. Only the values for the current time point are retained, because the process must be repeated each quarter as the new shocks are revealed. The simulations were conducted using the IDS or “stack-time” algorithm for forward-looking, non-linear models described by Armstrong, Black, Laxton, and Rose (1995). This algorithm is reasonably fast and very robust, both important features for the experiments that we consider in this paper.

The information structure is such that the monetary authority observes the shocks once they have been realized, but it does not know the future shocks. The monetary authority also knows the autocorrelation and cross-correlation structure of the shocks as well as the structure of the model, and it forms model-consistent expectations of inflation based on this information. The monetary authority in the model is, therefore, considerably better informed than real-world central banks, which face uncertainty both about the structure of the economy and about the precise shocks. In this regard, our results understate the inflation control problem.

The shocks are based on the estimated VAR as described in Section 2, with the exception of the rsl shocks. Recall that, in the general case of our VAR approach, there are two sources of variation in rsl. There are exogenous shocks to rsl (which produce the impulse responses shown in the sixth column of Figure 1) and there are shocks to rsl that are induced by exogenous shocks to the other five variables in the VAR so that the model produces the VAR response of rsl (as depicted in the bottom row of Figure 1) over the first five quarters. To examine the stochastic behaviour of the economy under the CPAM reaction function, we turn off all the induced shocks to the rsl equation that arise as a result of shocks to the other five variables in the VAR. Providing this response is the role of the reaction function in the model. For the exogenous shocks to rsl, we consider two cases. In the first case, the exogenous rsl shocks are also set to zero so the only source of variation in rsl is from the monetary response to shocks to the other five variables in the VAR. In the second case, we include exogenous monetary shocks but reduce their standard deviation to half of the historical estimate.

The rationale for setting the exogenous rsl shocks to zero is that these reflect discretion on the part of the monetary authority, whereas under the contingent rule embodied in the reaction function there would be no discretion. The impulse responses to the rsl shock in the VAR (shown in the last column of Figure 1) do indeed depict the textbook effects of a discretionary monetary contraction. In our view, however, eliminating the exogenous rsl shocks entirely is likely to understate the control problem facing the monetary authority. The VAR controls for the effects of the variables in the system on rsl, leaving the effects of other determinants of rsl to be captured by the residual in the rsl equation. One of these determinants is discretionary monetary policy, but the innovations in the rsl equation undoubtedly also reflect other influences, such as movements in U.S. interest rates, changes in risk premiums in the term structure, and monetary policy "errors" attributable to the Bank of Canada's having considerably less information than the monetary authority in the model. While following a monetary policy rule would eliminate discretionary monetary shocks, these other exogenous sources of fluctuations in rsl are likely to be more enduring. This suggests that some component of the rsl shock should be included in the stochastic simulations. Our second set of simulations therefore includes an exogenous rsl shock that has the same dynamic properties as estimated in the VAR, but with a standard deviation of half the VAR estimate. We view this "half" assumption as providing a reasonable upper bound on the proportion of rsl shocks that reflect non-monetary factors.

The results for both sets of simulations are reported in Table 2. The table reports the standard deviations and the AR(1) coefficients for selected variables in the model. The standard errors associated with these moments are reported in parentheses. As a basis of comparison for the model simulations, the first column of the table reports comparable statistics for the variables in question over the period 1961Q1 to 1992Q4, the estimation period for the VAR. With the exception of the yield spread, all the variables are logged. To the extent possible, non-stationary variables are defined in ratio form to induce stationarity. This was done to avoid the common practice of HP filtering non-stationary variables before computing moments, since it has recently been shown that this practice can induce spurious correlations (Cogley and Nason 1995). Thus, for example, the components of demand are expressed as ratios of output.[22] The exception is output itself, which is measured in terms of its quarterly growth rate at annual rates (Δy), its four-quarter growth rate ($\Delta_4 y$), and the deviations from its HP trend $(y - y^{HP})$. The measure of inflation we focus on is the four-quarter rate of change of the consumer price index (CPI) (πc_4), since it is this measure that the monetary authority aims to control in the model.

Looking first at the results with no rsl shocks, we see that standard deviations and the AR(1) coefficients for the various measures of detrended output in the model are very similar to those in the data. In addition, the main demand components display similar behaviour in the model relative to

22. Historically, some of the components of demand as a ratio of output do exhibit trends. Imports and exports, for example, have been growing as a proportion of output as a result of ongoing trade liberalization. In the model, of course, these ratios are all stationary in the stochastic steady state, since the model abstracts from institutional changes such as trade liberalization. To render the historical data comparable to the model data, we detrend the demand component ratios using a linear-quadratic time trend. There is, therefore, a sense in which the data are prefiltered, but this filtering problem is a second-order one compared with applying a business-cycle filter to the demand component itself.

Table 2

The Stochastic Behaviour of the Economy Using the Base-Case CPAM Rule

			Model			
	History		**No *rsl* shock**		**With *rsl* shocks**	
Variable	**Standard deviation**	**AR(1) coefficient**	**Standard deviation**	**AR(1) coefficient**	**Standard deviation**	**AR(1) coefficient**
Δy	4.1	0.33	4.0 (0.5)	0.37 (0.14)	4.1 (0.4)	0.47 (0.11)
$\Delta_4 y$	2.7	0.87	2.7 (0.5)	0.86 (0.03)	3.0 (0.5)	0.88 (0.04)
$y - y^{HP}$	1.6	0.84	1.8 (0.4)	0.87 (0.04)	2.1 (0.4)	0.89 (0.04)
$(c + i) / y$	1.5	0.86	1.6 (0.2)	0.89 (0.02)	1.6 (0.2)	0.89 (0.02)
c / y	1.6	0.87	0.9 (0.1)	0.95 (0.01)	1.0 (0.2)	0.95 (0.01)
i / y	1.6	0.93	1.0 (0.1)	0.86 (0.03)	1.0 (0.1)	0.86 (0.03)
$tbal / y$	1.9	0.87	1.6 (0.2)	0.93 (0.02)	1.8 (0.2)	0.93 (0.02)
g / y	1.3	0.92	0.5 (0.2)	0.95 (0.02)	0.7 (0.2)	0.95 (0.02)
rsl	1.2	0.77	1.5 (0.2)	0.88 (0.05)	1.8 (0.3)	0.84 (0.04)
rn	3.5	0.94	2.8 (0.4)	0.90 (0.04)	3.3 (0.6)	0.87 (0.04)
z	6.7	0.95	4.1 (0.5)	0.92 (0.02)	4.7 (0.8)	0.92 (0.02)
πc_4	3.1	0.97	0.8 (0.1)	0.94 (0.02)	1.1 (0.2)	0.95 (0.01)

Note: Standard errors are reported in parentheses.

history; this is particularly the case for $(c + i)/y$ and for the $tbal/y$.[23] In contrast, the behaviour of inflation in the model is quite different from history. The standard deviation of inflation in the model is only 0.8 percentage points as compared with 3.1 percentage points in the historical data. This improved inflation control translates into less variability in both the short-term nominal interest rate (*rn*) and the real exchange rate (*z*) relative to history, despite a slightly larger standard deviation for *rsl* in the model as compared with history.

The marked improvement in inflation control relative to history is consistent with the fact that the objective of monetary policy in the model is to return inflation to its target of two per cent, whereas over much of the historical sample, monetary policy was not geared towards an inflation target, and inflation varied considerably from its historical mean.

23. The properties of the components of $c + i$ in the model do not match those in the data as closely as do those for the total. This reflects that fact that in splitting the $c + i$ shocks from the VAR in order to put them in the model, we assume that the consumption and investment shocks are perfectly correlated. The standard deviation of g/y is also lower in the model than in the data since there is no fiscal shock in our analysis.

Interestingly, the degree of inflation control in the model fits well with the experience since 1992, when monetary policy in Canada has been guided by explicit inflation targets. In fact, the standard deviation of CPI inflation (and inflation in the CPI excluding food and energy) relative to the midpoint of the inflation control range was 0.9 percentage points over the 1992 to 1996 period.

Including an exogenous *rsl* shock in the stochastic simulations has the effect of slightly increasing the volatility of most variables in the model, but overall the results are similar. For example, the standard deviation of the four-quarter growth rate of output rises from 2.7 percentage points with no *rsl* shocks to 3.0 percentage points with *rsl* shocks, and the standard deviation of inflation rises from 0.8 percentage points to 1.1 percentage points.

Table 3 reports some additional information on the degree of inflation control using the CPAM reaction function. Note, first, that in simulations both with and without the *rsl* shocks, this reaction function does succeed in delivering an average rate of inflation that is not significantly different from its target. In both cases, the mean inflation rate is less than half of one standard error from the target. This suggests that, despite the asymmetry in the Phillips curve in the model, a symmetric reaction function is sufficient to achieve the inflation target on average.

The variability of inflation around its target is conveniently measured by the proportion of time that inflation is outside a given band around the target. For example, with no *rsl* shocks, inflation is predicted to be outside bands of plus or minus (+/–) 1 percentage point 27 per cent of the time, and the median duration of departures from the bands is 4.7 quarters. Increasing the bands to +/– 1.5 percentage points reduces the proportion of time outside the bands to only 9 per cent; the 95 per cent confidence band for inflation around its target is +/– 1.9 percentage points. With the *rsl* shocks included, the comparable bands all increase slightly. In particular, inflation is now outside bands of +/– 1 percentage point 38 per cent of the time, and the 95 per cent confidence interval around the target is +/– 2 percentage points.

These bands are relatively narrow compared with those suggested in previous studies for the United States, Canada, Australia, and the United Kingdom. Taylor (1979) and Fuhrer (1994) estimate a 95 per cent confidence interval of about +/– 3.5 percentage points for the United States, and Fillion and Tetlow (1994) and Crawford and Kasumovich (1996) reach quantitatively similar conclusions for Canada. Stevens and Debelle (1995) conclude that the 95 per cent confidence interval for inflation in Australia is between +/– 2.4 and +/– 3.6 percentage points. For the United Kingdom, Haldane and Salmon (1995) find considerably wider bands, ranging between

Table 3

Inflation and Output Variability with the Base-Case CPAM Rule

	No *rsl* shocks	With *rsl* shocks	With *rsl* shocks but no cross-correlations
RMSD($\pi c_4 - \pi^*$)	0.85	1.06	1.60
Standard deviation (πc_4)	0.85 (0.20)	1.05 (0.20)	1.60 (0.23)
MEAN($\pi c_4 - \pi^*$)	0.02 (0.20)	–0.11 (0.32)	–0.02 (0.59)
RMSD(*ygap*)	2.87	3.67	3.26
Standard deviation (*ygap*)	2.67 (0.67)	3.33 (0.77)	2.99 (0.68)
MEAN(*ygap*)	–0.90 (0.67)	–1.45 (0.73)	–1.17 (0.82)
Percent of time outside the bands			
+/– 1.0	27	38	58
+/– 1.5	9	17	40
+/– 2.0	3	5	25
+/– 2.5	0	0	13
Median duration in quarters of departures from bands			
+/– 1.0	4.7	5.9	7.0
+/– 1.5	4.2	4.1	6.8
+/– 2.0	2.2	3.0	5.2
+/– 2.5	0.0	0.5	4.5

Notes: RMSD is the root-mean-squared deviation. The mean-squared deviation (MSD) for inflation is the variance of inflation plus the squared deviation of inflation from target. The RMSD is the square root of the MSD. As for the other moments, the reported RMSD is calculated by computing the RMSD for each replication of 128 periods, and then averaging over the 10 replications. Standard errors are reported in parentheses.

+/– 7 and +/– 9.5 percentage points. Our results are considerably more encouraging for the prospects for inflation control.[24]

Table 3 also compares the economies with and without *rsl* shocks in terms of their output performance. Introducing the *rsl* shock in the simulations increases the standard deviation of the output gap from 2.7 per cent to 3.3 per cent. Because of the asymmetry in the Phillips curve in the model, the variability of output also affects the average level of output

24. For New Zealand, Turner (1996) concludes that the 95 per cent confidence interval is a little over +/– 1.5 percentage points. This is actually a slightly narrower band than we obtain for Canada. However, Turner's inflation measure is the Reserve Bank of New Zealand's measure of underlying or core inflation, and this is considerably less volatile than the total CPI, which is the price concept we are using.

in the economy relative to the deterministic measure of potential.[25] Both with and without the *rsl* shocks, the mean of the output gap is negative, but the absolute size of this negative mean shift in output is larger when output is more variable. The negative mean shift arises because inflation is more sensitive to excess demand than to excess supply, so the cumulated excess supply gaps in the model (when output is measured against the deterministic measure of potential in the model) have to be larger than the cumulated excess demand gaps to prevent inflation from drifting upwards over time. In the simulations with the *rsl* shocks, the average level of output that can be achieved over time is 1.5 per cent below the deterministic level of potential output (that is, the average level of output that could be achieved over time if there were no variability in output). With no *rsl* shocks, the mean shift in output is cut approximately in half.

As a final word on Table 3, note that the third column reports the results when the *rsl* shock is included, but the shocks to all six variables are restricted to be independent from one another. That is, shocks are autocorrelated (based on an AR(1) process estimated from the VAR impulse responses), but not cross-correlated.[26] This experiment is presented to show the impact of allowing for cross-correlated shocks in our base-case results. As reported, when the shocks are restricted to be independent, the standard deviation of inflation is higher than with cross-correlated shocks, while for output the standard deviation is lower. This suggests that there is some positive correlation between aggregate demand shocks and aggregate supply shocks. To the extent that the cross-correlations among the shocks as derived from our VAR approach do not reflect deep structure, the results in columns one and two may understate the inflation control problem on the margin.

4 The Prospects and Implications of Improved Inflation Control

4.1 Reaction functions in the EF-I class

The results presented above suggest that the base-case CPAM reaction function is a reasonable rule in that it does a good job of controlling inflation (indeed, very good when measured against the experience between 1961 and 1992), while at the same time producing volatility measures for other variables, such as output, interest rates, and the real exchange rate, that

25. For an analysis of this issue, see Laxton, Rose, and Tetlow (1993b). See also Laxton, Meredith, and Rose (1995); Clark, Laxton, and Rose (1996, 1997); Debelle and Laxton (1996); and Clark and Laxton (1997).

26. The standard deviations of the shocks are scaled so that they are the same in the second and third columns of Table 3.

are in line with or lower than historical values. The base-case CPAM rule, however, reflects only one possible parameterization of rules in the EF-I class, and there may be alternative parameterizations that produce more desirable outcomes. In this section, we explore that possibility.

In particular, we consider reaction functions in the EF-I class and examine whether there are rules in this class that outperform the base-case CPAM rule in the sense that they reduce both the root-mean-squared deviation of inflation from its target and the root-mean-squared deviation of output from the deterministic level of potential output. We use the root-mean-squared deviation (RMSD) instead of simply the standard deviation in an effort to combine information about the performance of alternative reaction functions in terms of the average levels of inflation and output they produce, as well as the variability around these averages. So, for example, the RMSD will penalize a reaction function for failing to hit the inflation target on average, as well as for variability in inflation. The use of the RMSD is particularly important for the output gap, since lowering its standard deviation will also tend to reduce the mean shift in the output gap that arises because of the asymmetry in price adjustment. Since economic welfare depends on the average level of output as well as the variability, we view the RMSD as a more appropriate yardstick. Nevertheless, we realize that it is arbitrary to weight the level and variability effects equally (and it likely underweights the level effects relative to their welfare implications).

To keep the problem computationally manageable, we restrict our attention to rules in the EF-I class that have the same form as the CPAM rule in the sense that they use a two-period moving average of the expected deviation of inflation from the target. Specifically, we conduct stochastic simulations with the model using reaction functions of the form

$$\begin{aligned} rsl = rsl^* + \theta[1.6\bullet(\pi c_4(j) - \pi^*) \\ + 1.8\bullet(\pi c_4(j+1) - \pi^*)] + \xi^{rsl}, \end{aligned} \tag{6}$$

and vary θ, the weight on expected deviations of inflation from its target (as a multiple of the base-case weights), and j, the number of quarters into the future the monetary authority looks in evaluating the expected deviation of inflation from its target. The CPAM rule is the special case with $\theta = 1$ and $j = 7$. Simulations were performed for reaction functions with $j = 0, 1, 2, \ldots, 10$; and for each j, 26 different values of θ were considered, ranging from 0.25 to 20. Stochastic simulations were therefore conducted for 286 different reaction functions in the EF-I class. To make this computationally feasible, we reduced the number of replications from 10 to 4, and the length of the simulations from 136 periods to 64 periods (again dropping the first 8 periods). This still implies

Figure 2

Efficient Policy Frontier for Reaction Functions in the EF-I Class in Inflation-Output Space

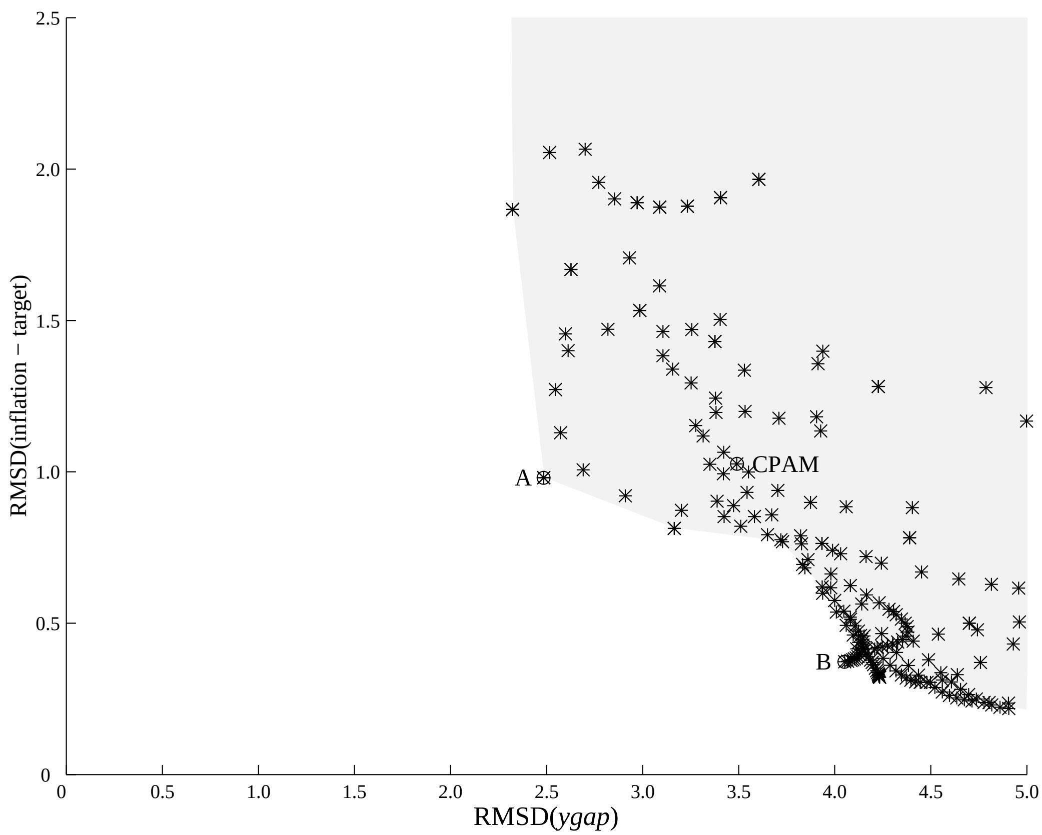

(286 x 4 x 64) = 73,216 simulations.[27] All the experiments were conducted using the shocks for the non-monetary variables as derived from the VAR, together with an exogenous *rsl* shock with half its historical standard deviation (as with "*rsl* shocks included" in Tables 2 and 3).

The results of this exercise are summarized in Figure 2. The figure plots the RMSD of the deviation of inflation from its target against the

27. This is a large-scale computational problem. Our Sun Sparc Ultra compute server would have to run continuously for about five days to perform all these simulations. Fortunately, we were able to divide the problem up across 26 Sparc Ultra workstations running simultaneously. For each j, the 26 values of θ were allocated among the 26 machines. Using this method the 73,216 simulations ran in about 10 hours.

RMSD of the output gap for each of the 286 reaction functions considered.[28] The hull of these points is shaded to provide a general impression of the space spanned by reaction functions in this class. Along any ray from the origin, points closer to the origin are "better" in the sense that they produce lower RMSDs of both inflation and output gaps. The points closest to the origin are said to be on the "efficient policy frontier" for this class of reaction function since, for each of these points, there is no other point that has a lower RMSD for both variables.

The data point in the figure with the label "CPAM" is the outcome using the base-case CPAM reaction function. Note that this point is not on the efficient policy frontier; there are many points below and to the left of the CPAM point. These points all share the feature that the weight in the reaction function on expected deviations of inflation from its target is larger than in the CPAM rule (that is, $\theta > 1$). Thus, by reacting more vigorously to deviations of inflation from its target, the monetary authority can reduce the RMSD of both variables. This likely arises because, with demand shocks, inflation is controlled by leaning against the output gap. The fact that the CPAM base-case reaction function is not on the efficient policy frontier suggests that the demand shocks emanating from our estimated VAR have sufficiently important and persistent effects that a more vigorous rule decreases the variability of both inflation and output.

Along the efficient policy frontier for reaction functions in the EF-I class, there is generally a trade-off between the RMSDs of the inflation and output gaps of the kind suggested by Taylor (1979, 1994), although the trade-off we compute is not always convex to the origin. The trade-off arises for at least two reasons. First, as stressed by Taylor, while controlling inflation in the face of demand shocks requires leaning against the output gap, with price shocks (that is, shifts in the expectations-augmented Phillips curve), returning inflation to its target will tend to widen the output gap in the short run, thereby producing a trade-off between the volatilities of output and inflation. In dynamic models where targets cannot be hit contemporaneously, there is also a second reason to expect a trade-off. Even if there are only demand shocks, as the monetary authority pursues increasingly vigorous reaction, there will come a point when controlling inflation more tightly in the short run will begin to induce larger secondary cycles in output. Thus, while a more vigorous reaction function may succeed in doing a better job of

28. There are in fact a little fewer than 286 points on the figure because some points are off the scale of the figure, and because, for some reaction functions, the simulations failed to converge for some draws, in which case we dropped the associated point. Failures actually happened relatively infrequently; in cases when they did occur, we interpreted this as evidence that the associated reaction function could not produce acceptable outcomes.

Table 4

Comparing Some Alternative Reaction Functions in the EF-I Family

	Base-case CPAM j = 7, θ = 1.0	**Point A** j = 8, θ = 3.5	**Point B** j = 4, θ = 20
RMSD($\pi c_4 - \pi^*$)	1.06	1.16	0.40
Standard deviation (πc_4)	1.05 (0.20)	1.14 (0.19)	0.40 (0.05)
MEAN($\pi c_4 - \pi^*$)	–0.11 (0.32)	–0.18 (1.08)	–0.01 (0.11)
RMSD(*ygap*)	3.67	2.76	4.07
Standard deviation (*ygap*)	3.33 (0.77)	2.46 (0.46)	3.80 (0.78)
MEAN(*ygap*)	–1.45 (0.73)	–0.77 (1.13)	–1.30 (0.90)
Standard deviation (*rsl*)	1.80 (0.27)	2.56 (0.38)	4.26 (0.33)
Standard deviation (*rn*)	3.30 (0.61)	4.64 (0.84)	6.69 (0.48)
Standard deviation (*z*)	4.66 (0.84)	3.51 (0.66)	4.98 (0.88)
Per cent of time outside the bands			
+/– 1.0	38	47	1
+/– 1.5	17	30	0
+/– 2.0	5	18	0
Median duration in quarters of departures from bands			
+/– 1.0	5.9	10.6	1.5
+/– 1.5	4.1	9.1	0.0
+/– 2.0	3.0	6.2	0.0

Note: RMSD is the root-mean-squared deviation. The mean-squared deviation (MSD) for inflation is the variance of inflation plus the squared deviation of inflation from target. The RMSD is the square root of the MSD. As for the other moments, the reported RMSD is calculated by computing the RMSD for each replication of 128 periods, and then averaging over the 10 replications. Standard errors are reported in parentheses.

stabilizing inflation, the secondary cycles in output associated with this policy will begin to destabilize output, resulting in a trade-off.

The points labelled A and B in Figure 2 provide two points on the efficient policy frontier for the EF-I rules that are at quite different ends of this trade-off. Interestingly, these two points are also associated with quite different reaction functions. The reaction function at point A has $j = 8$ and $\theta = 3.5$, while at point B, $j = 4$ and $\theta = 20$.[29] Table 4 reports more

29. Note that, since point B has the highest value of θ we considered, we also ran some experiments with slightly higher values of θ with $j = 4$, but found that the results were essentially the same as for point B. Indeed, as suggested by the density of points in the vicinity of B, there is a relatively broad range of reaction functions that produce similar results. These reaction functions all share the feature that j is low (2 to 4). In general, points along the frontier that are to the right of point B involve shorter leads than point B, while to the left of point B (up to point A), they involve longer leads (from 5 to 8).

detailed information on the stochastic behaviour of the economy under the reaction functions at points A and B, as well for the CPAM rule. To put points A and B on a comparable basis with the CPAM rule and to increase the precision of the calculations, we re-ran the simulations for these points using 10 replications over the longer horizon of 128 usable periods.[30]

Comparing the point-A reaction function first with the CPAM rule reveals that these two rules deliver similar outcomes in terms of inflation control. The advantage of the point-A rule relative to the CPAM rule comes in the variability of the output gap. Moving from the CPAM rule to the point-A rule reduces the standard deviation of the output gap from 3.3 per cent to 2.5 per cent, and this cuts the mean shift in the level of output approximately in half, from −1.5 per cent to −0.8 per cent of potential. This is an economically significant improvement. The more vigorous point-A reaction function does produce more variability in interest rates—the standard deviation of rsl rises from 1.8 percentage points with the CPAM rule to 2.6 percentage points at point A, and for the rn the increase is from 3.3 to 4.6 percentage points. As reported in Table 2, the standard deviation of *rn* over history is 3.5, so nominal rates at point A are quite variable. This suggests that it may be possible to view the CPAM rule as an efficient rule if interest rate variability is viewed as undesirable. Interestingly, however, the variability of the real exchange rate is lower using the point-A rule, so the case for the CPAM rule requires interest rate volatility to be more harmful than exchange rate volatility.

Moving from point A to point B involves trading off more variability in the output gap against better inflation control. At point B, the standard deviation of inflation is reduced to only 0.4 percentage points, implying that inflation would be outside bands of +/– 1 percentage point *only 1 per cent of the time*. This dramatic improvement in inflation performance comes, however, at considerable cost in terms of output variability relative to point A (and to a lesser extent relative to the CPAM rule). In particular, in moving from A to B, the standard deviation of the output gap rises from 2.5 percentage points to 3.8 percentage points.

The very different outcomes at points A and B reflect the different channels through which inflation is controlled at these points. The reaction function at point A (like the CPAM rule) has a relatively long lead on the deviation of inflation from the target. Inflation is controlled over this horizon principally through the effect of changes in short-term interest rates and the exchange rate on the output gap, together with the influence of the output

30. Therefore, the results in Table 4 do not match those in Figure 2 exactly. However, they are reasonably close, which gives us some confidence that the smaller runs of four replications with 56 usable periods are not subject to critical small-sample problems.

gap on inflation. Since the lags in this transmission mechanism are about two years, inflation control is somewhat imprecise, but because it is achieved by dampening cycles in output, output variability is relatively low. In contrast, the lead on the deviation of inflation from the target in the point-B rule is quite short, at only four to five quarters. Over this horizon, monetary policy has little influence on inflation through the output gap, but given the significant import share of consumer goods, it can affect inflation over this horizon through its effect on the exchange rate. For example, if inflation is expected to be above the target four to five quarters ahead, the monetary authority would raise interest rates, triggering an appreciation of the exchange rate. This appreciation feeds through to lower consumer prices over the following year or so, thereby bringing inflation back to the target.[31] However, for all the values of θ we considered, rules with short leads (0 to 5) produce considerably more output volatility than rules with leads that correspond to the monetary control lag in the model through the output gap (leads 7 to 9). With high θ's, the short-lead rules are on the frontier, since they produce lower inflation variability than the rules with longer leads. But not surprisingly, these high-θ, short-lead rules produce considerable variability in interest rates and the exchange rate, in addition to output.

This high volatility of nominal interest rates at point B also illustrates a potential problem with the feasibility of the point-B rule. With a 2 per cent inflation target and an equilibrium real interest rate of about 3 per cent, the average nominal interest rate is about 5 per cent. At point B, one standard deviation of the nominal interest rate is 6.7 percentage points, so the nominal rate is frequently going negative in our simulations. Since money does not enter the model explicitly, a negative nominal rate poses no special problems in the model. In a richer model, however, in which money demand becomes unbounded (or at least very large) as the nominal interest rate approaches zero, the nominal interest rate would not go negative. This would no doubt reduce the degree of inflation control that could in fact be

31. This exchange rate approach to inflation control has been emphasized by the Reserve Bank of New Zealand. Grimes and Wong describe the Reserve Bank's monetary framework as follows: "Policy is aimed explicitly at the objective of maintaining price stability, and no formal intermediate targets are adopted. The exchange rate is used, however, as a calibration device to determine appropriate monetary settings (the level of bank reserves) to deliver the inflation target given other forecast inflation influences" (Grimes and Wong 1994, 195). It is also clear that the focus on the exchange rate reflects its influence on prices: "the exchange rate acts as the principal guide for calibrating how much monetary settings should be changed when forecast inflation deviates from the target range. This is not because significance is attached to the exchange rate level per se, but because it is considered to be a strong influence on the future price level and it has the additional advantage that it can be monitored continuously" (Grimes and Wong 1994, 181-82).

achieved using the point-B reaction function. It also shows the potential policy relevance of the so-called Summers effect (Summers 1991) if policymakers both target a low rate of inflation and attempt to keep inflation within a very narrow band of this target.[32]

A final noteworthy feature of the results in Figure 2 is that the surface of points between A and B is concave to the origin. This concavity would seem to reflect the fact that the control lag to inflation through exchange rate pass-through is about one year (which is the basis of the point-B rule), and through the output gap it is about two years (which is the basis of the point-A rule), but between these horizons, inflation control is less effective.[33]

4.2 Reaction functions in the C-IY class

Policy rules in the EF-I class are only one set among many. The literature on monetary policy rules is large and many types of rules have been studied.[34] Among these, the nominal income targeting rule has been frequently advocated as providing a reasonable balance between concern for output and prices (see, for example, Meade 1978, Tobin 1980, and McCallum 1984). More recently, Taylor (1993, 1994) has popularized a version of the nominal income rule in which the real interest rate reacts with equal coefficients to the contemporaneous values of inflation relative to the target and the output gap. The Taylor rule and generalizations that allow for different weights on contemporaneous values of inflation and the output gap (which we call C-IY rules) have attracted considerable attention, particularly among researchers at central banks (see Judd and Motley 1991, Fuhrer 1994, Haldane and Salmon 1995, Stevens and Debelle 1995, Levin 1996, and Tetlow and von zur Muehlen 1996).

On the theoretical front, recent work by Svensson (1996b) and Ball (1997) compares EF-I and C-IY rules. They both demonstrate that rules in the EF-I class are analytically equivalent to rules in the C-IY class in a small, *linear*, closed economy model. The reason is that in their simple models, the current output gap and the current rate of inflation are efficient predictors of future inflation. Thus, even if the only concern is to minimize

---

32. Black, Coletti, and Monnier (this volume) find the Summers effect is rarely binding and consequently has minor welfare implications with a 2 per cent inflation target using the base-case QPM policy rule (which is similar to the base-case CPAM rule). Our results suggest that the welfare implications of the Summers effect may be more important for some alternative policy rules, but we leave this issue for future research.

33. To see if we could obtain points along a straight line between A and B, we experimented with linear combinations of the point-A and point-B rules. We found that this produced points closer to the origin relative to the shaded area in Figure 2 between A and B, but the shape of the resulting frontier was still concave between A and B.

34. See McCallum (1997) for a recent survey.

the deviation of expected inflation from the target, when expressed in terms of contemporaneous variables the policy rule will include a weight on the current output gap.

In CPAM, however, rules in the C-IY and EF-I classes are not equivalent in this sense, since the model is non-linear and the open economy dimension leads to inflation being affected importantly by the exchange rate channel. Moreover, in the context of a small-scale, closed economy model, Clark, Laxton, and Rose (1996, 1997) have shown that the performance of rules that are based on contemporaneous values can be quite different in linear and non-linear models. Comparing rules in the EF-I and C-IY classes is therefore an interesting exercise.

To compare these rules in the open economy, non-linear framework of CPAM, we computed the efficient policy frontier for rules of the form

$$rsl = rsl^* + \delta(\pi c_4 - \pi^*) + \lambda(ygap) . \tag{7}$$

These rules have the same general form as the Taylor rule, with an important difference. In the Taylor rule, the controller is the real interest rate (rather than the yield spread), measured as the nominal rate less the most recent four-quarter inflation rate. For the United States, Taylor (1993) suggests that his rule with weights of 0.5 on both inflation and the output gap provides a good representation of the Federal Reserve's historical rule.

We examine the efficient policy frontier for rules of the form of equation (7) by varying the weights on inflation (δ) and the output gap (λ). Specifically, we consider 10 different values for λ between 0.5 and 10, and 26 values for δ between 0.25 and 20. As above, for the full frontier, the computed moments for each set of parameters are based on 4 replications of samples of 56 quarters, and we repeat the simulations for selected points of interest using 10 replications of samples of 128 quarters. Figure 3 depicts the efficient policy frontier for the rules in the C-IY class, together with the frontier already obtained for rules in the EF-I class, and Table 5 reports more accurate and more detailed results for selected C-IY rules of interest. Several features of the results are noteworthy.

First, the C-IY rules that lie along the frontier all have high coefficients relative to the Taylor weights of 0.5 and 0.5.[35] For example, the weights corresponding to the points marked "C," "D," and "E" in Figure 3 range from 9.0 to 12.0 on inflation and from 2.0 to 6.0 on output. In fact, with the Taylor weights, the C-IY rule has an RMSD for inflation of 3.3 percentage points, putting it over the top of Figure 3. In models of the

35. Recall, however, that our C-IY rules use yield spread as the monetary variable, rather than the real interest rate, as in the original Taylor formulation.

Figure 3

Efficient Policy Frontiers for Reaction Functions in the C-IY and EF-I Classes in Inflation-Output Space

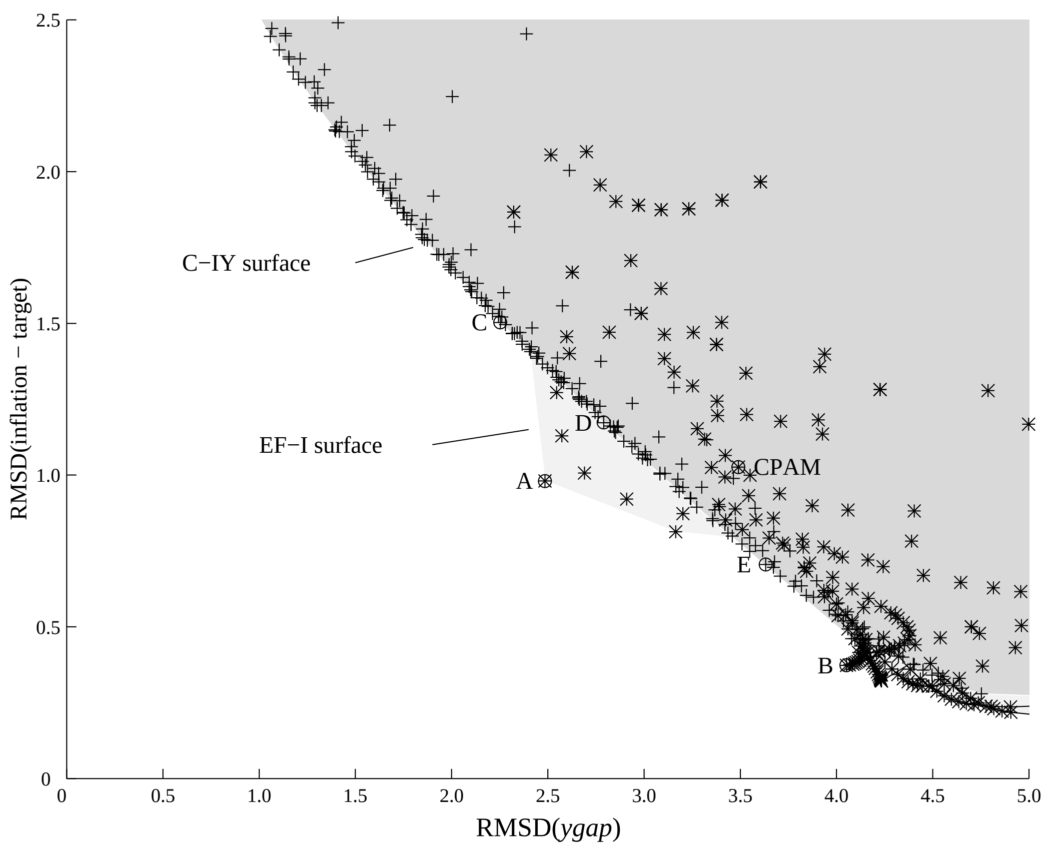

U.S. economy, Fuhrer (1994), Levin (1996), Tetlow and von zur Muehlen (1996), and Clark, Laxton, and Rose (1997) also find that efficient rules in the Taylor family have considerably higher coefficients than Taylor's.

Second, note that the C-IY rule spans a greater range of outcomes in inflation-output RMSD space, but over the range that is also spanned by the EF-I rules, the efficient frontiers for both sets of rules are relatively similar. The EF-I rules appear to do slightly better in the vicinity of point A in Figure 3, but the difference is not large enough to be significant, given the stochastic nature of the exercise and the limited numbers of replications.[36]

36. This is evident comparing the moments for points A and D in Tables 4 and 5, respectively. In Figure 3 (based on 4 replications over 56 quarters), point A appears to dominate point D, but in Tables 4 and 5 (with 10 replications over 128 quarters) the two points are shown to be very similar in terms of the RMSDs of output and inflation.

Table 5

Comparing Some Alternative Reaction Functions in the C-IY Family

	Point C $\lambda = 6.0$, $\delta = 9.0$	**Point D** $\lambda = 4.0$, $\delta = 10.0$	**Point E** $\lambda = 2.0$, $\delta = 12.0$
RMSD($\pi c_4 - \pi^*$)	1.41	1.09	0.65
Standard deviation (πc_4)	1.32 (0.32)	1.01 (0.23)	0.61 (0.12)
Mean($\pi c_4 - \pi^*$)	0.50 (0.39)	0.39 (0.28)	0.24 (0.15)
RMSD(*ygap*)	2.28	2.78	3.60
Standard deviation (*ygap*)	2.15 (0.46)	2.61 (0.55)	3.37 (0.70)
Mean(*ygap*)	–0.64 (0.53)	–0.84 (0.62)	–1.16 (0.75)
Standard deviation (*rsl*)	4.23 (0.31)	3.61 (0.36)	3.70 (0.41)
Standard deviation (*rn*)	6.64 (0.44)	5.70 (0.54)	5.78 (0.57)
Standard deviation (*z*)	4.00 (0.58)	4.30 (0.62)	4.87 (0.71)
Per cent of time outside the bands			
+/– 1.0	40	32	11
+/– 1.5	27	16	2
+/– 2.0	16	6	0
Median duration in quarters of departures from bands			
+/– 1.0	8.3	6.9	4.2
+/– 1.5	7.2	6.2	2.2
+/– 2.0	7.1	5.8	0.0

Notes: RMSD is the root-mean-squared deviation. The mean-squared deviation (MSD) for inflation is the variance of inflation plus the squared deviation of inflation from target. The RMSD is the square root of the MSD. As for the other moments, the reported RMSD is calculated by computing the RMSD for each replication of 128 periods, and then averaging over the 10 replications. Standard errors are reported in parentheses.

The results in Tables 4 and 5 reveal that this similarity of the C-IY and EF-I rules in inflation-output-RMSD space hides some other important differences. Note, in particular, that while rules in the EF-I class generally hit the inflation target on average, rules in the C-IY class produce a systematic positive inflation bias relative to the target. At point D, for example, the mean inflation rate is 0.4 percentage points above the target.[37] This bias arises because of the asymmetry in the Phillips curve. With a

37. This bias of 0.4 percentage points is 1.5 standard deviations from zero, so it is not quite significant at the 95 per cent level given the precision of our stochastic analysis. However, this positive inflation bias was a common feature across the different calibrations of the C-IY rules, and for calibrations that produce points off the frontier, the bias is often much larger. For example, with the Taylor weights of 0.5 on output and inflation, the average inflation rate is 2.5 percentage points above the target (with a standard deviation of 1.2 percentage points).

Figure 4

Efficient Policy Frontiers for Reaction Functions in the C-IY and EF-I Classes in Inflation-Interest Rate Space

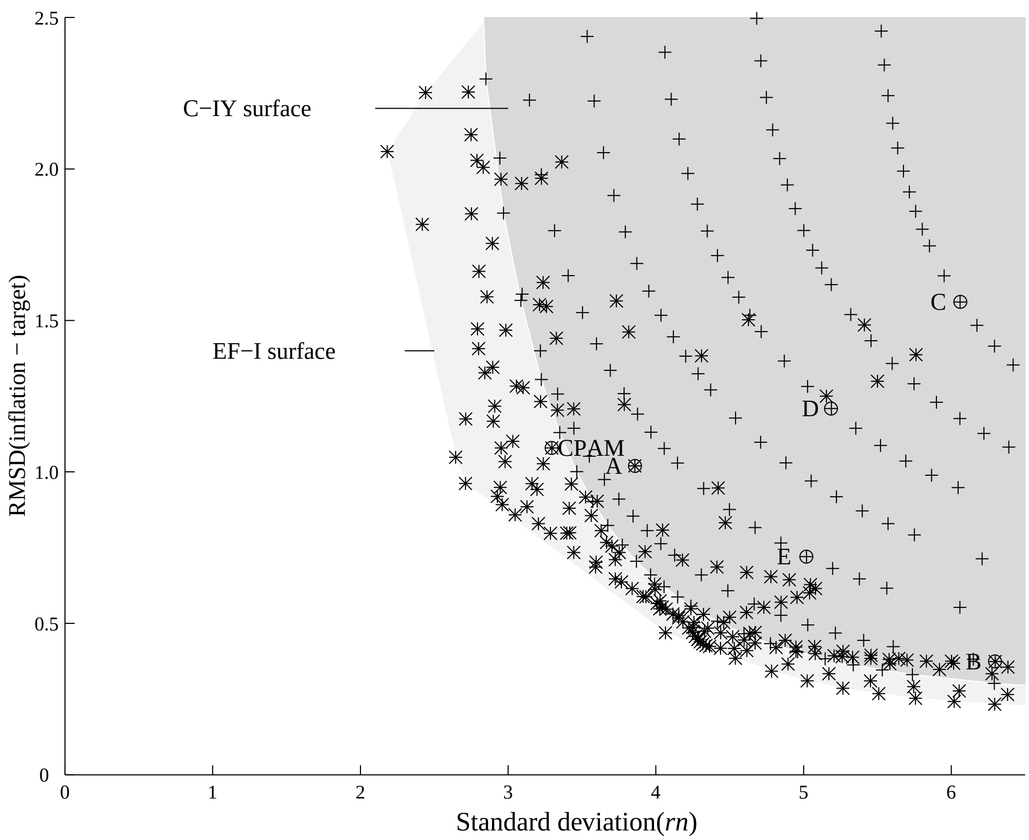

convex Phillips curve, the deterministic level of potential output $y*$ is not attainable on average over time. The resulting negative mean shift of output relative to this deterministic level of potential output (which is –0.8 at point D) produces an average value of the monetary controller that is too low to achieve the inflation target on average.

The second important difference between the EF-I and C-IY rules is that for the latter, interest rates are considerably more volatile. For example, at point A on the EF-I frontier, the standard deviations of rsl and rn are 2.6 and 4.6 percentage points, respectively, as compared with 3.6 and 5.7 percentage points at D on the C-IY frontier. Figure 4 illustrates the greater interest rate volatility of C-IY rules more generally. It plots the RMSD of inflation relative to the target against the standard deviation of the nominal interest rate for the rules considered in the EF-I and C-IY classes. Note that in this dimension, the EF-I rules tend to dominate, with a relatively large number of points to the left of the frontier for the C-IY rules. Note also that the C-IY rules at points C, D, and E along the efficient frontier in inflation-output space are well back of the frontier in inflation-interest rate space, and well back of the EF-I point-A rule.

The higher variability of interest rates with the C-IY rules suggests that if there is any preference for low interest rate variability, the EF-I rules do considerably better. This conclusion is reinforced by the fact that, with more volatility in rates, negative interest rates are more frequent along the efficient frontier for C-IY rules. As a result, while points like A and D appear quite comparable in inflation-output space, once negative nominal rates are ruled out, point A will likely dominate.

The better performance of the EF-I rules, in terms of both their ability to hit the inflation target on average and the volatility of interest rates, reflects the fact that the EF-I rules are based on the model-consistent solution for inflation, whereas the C-IY rules use only contemporaneous information on inflation and output. The model-consistent solution takes account of the full structure of the model, as well as other useful information for inflation, such as exchange rate movements. In a non-linear, open economy model, this gives the EF-I rules an informational advantage that allows them to outperform the C-IY rules.

5 Price-Level Targeting, EF-IP Rules

We now turn to the issue of what the stochastic behaviour of the economy might look like if some element of price-level control were introduced into monetary policy. We do not consider a rule where the monetary authority acts to anchor the price level to some fixed value. Rather we consider the imposition of trend stationarity, where any drift in the price level away from a path consistent with the target rate of inflation is resisted and eventually eliminated. Thus, when a shock causes inflation to deviate from its target, the implications for the price level must be unwound, implying that inflation will not return monotonically to its target. For example, if a shock pushes inflation and the price level above their targets, inflation will have to fall temporarily below its target to bring the price level back to its target path.

A common argument is that, if it is costly to fight against inflationary shocks, then adding another restriction on the required outcome for the price level will make it costlier still. Thus, relative to an inflation-targeting rule, a price-level rule will imply more variability in output (see Lebow, Roberts, and Stockton 1992; Fillion and Tetlow 1994; and Haldane and Salmon 1995). In this setting, the case for price-level targeting rests on the benefits of reducing low-frequency price-level uncertainty, which must be traded off against the costs of greater output uncertainty.

As emphasized in Duguay (1994), Svensson (1996a), and Coulombe (this volume), however, this standard argument does not allow for possible changes in the way expectations are formed in response to the implied change in the properties of prices and inflation. In particular, if there were no unit root in the price level, the whole character of expectations formation would necessarily change. We know, for example, that a minimum mean-square error forecast from an autoregressive representation of the price level would have weights that sum to less than 1 (Sargent 1971). With an element of price-level control, there would be an expectation of reversion to the stationary mean of the price level, perhaps not monotonically, but enough to make expectations of *inflation* rates change sign relatively quickly. Contrast this with what happens in a world where inflation is stationary, and the log of the price level has a unit root. There would be mean reversion in this case too, but to the inflation rate. This would leave large extrapolative elements in expectations about the price level. In the EF-I world, where the drift is not controlled, we normally expect to see persistence in terms of inflation, and this would be reflected in expectations.

To take a preliminary look at whether price-level conditions might have a place as part of a policy of price stability, we consider reaction functions of the form:

$$\begin{aligned} rsl \; = \; & rsl^* + \theta[1.6 \bullet (\pi c_4(8) - \pi^*) + 1.8 \bullet (\pi c_4(9) - \pi^*)] \\ & + \tau \bullet pcgap(f), \end{aligned} \tag{8}$$

where $pcgap(f)$ is the expected difference between the f-quarter-ahead level of consumer prices and the price-level target path, where the price-level target path is defined as the level of consumer prices that is implied by a constant rate of inflation of π^* from some given starting point. Note that this reaction function reduces to a pure inflation-targeting rule with τ, the weight on the price-level gap, set to zero. Thus, the effects of adding an element of price-level control are considered on the margin. Note also that we have fixed the lead in the inflation component of the rule to be eight to nine quarters, which corresponds to the rule at point A in the class of EF-I rules. Our results with the EF-IP rules are therefore only partial; a more complete analysis would also allow the lead on the inflation component to vary.

As a first, and admittedly arbitrary, way of adjusting expectations to reflect the impact that the introduction of a price-level controller might have, we modify the equations for expected inflation as follows:

$$\begin{aligned} \pi^e \; = \; & 0.35 \bullet (0.7 \bullet \pi(-1) + 0.3 \bullet \pi(-2)) + 0.55 \bullet \pi_4(4) \\ & + 0.1 \bullet \pi^{*e} - 0.1 \bullet \eta \bullet pcgap(0). \end{aligned} \tag{9}$$

The new equation is as before—that is, as in equation (4)—with the exception that we have added the price-level gap term with a weight of $0.1 \cdot \eta$. The effect of this term is to reduce expected inflation on the margin when the current price level is above its target, and increase expected inflation when the current price level is below its target. When the price level is at its target, $pcgap$ is zero, so this term has no impact on expected inflation. In general, the weight on the price-level term should be a function of τ, the weight the monetary authority places on the price level in its reaction function (see Coulombe, this volume), but in this preliminary investigation we fix the weight as follows. We interpret the parameter η as telling us how fast the level gap is expected to be run off. For this experiment, we set η at 0.125, which implies a perception that an eighth of the level gap would be eliminated in the next quarter, all else being equal, so a given level gap would be 90 per cent eliminated over four years. We then set the weight on $\eta \cdot pcgap$ to 0.1, which, in our view, is small. In this respect, we view the adjustment to expectations as relatively minor and quite plausible.

The EF-IP reaction functions given by equation (8) require values for three parameters: the weights θ and τ, and the lead f in the $pcgap$ term. We consider a range of values for these parameters in two steps. In the first step, we set θ to the point-A value of 3.5, and vary τ and f. With $\tau = 0$, the reaction function therefore reduces to the point-A rule considered above, so the results with alternative values of τ and f, shown in Figure 5, should be compared with point A. The thin lines in Figure 5 show the results as we vary the value of τ, the weight on $pcgap$ in the reaction function, and there is one line for each of f = 4, 5, 6, and 7. We found that with an element of price-level control, the four replications we used to generate the EF-I frontier were insufficient. The results for the EF-IP rules therefore come from 10 replications with the same simulation length (16 years) as the EF-I points.

There are two conclusions from this preliminary exercise. First, the EF-IP rules produce a large number of points that lie closer to the origin than do points along the efficient frontiers for the reaction functions in either the EF-I or C-IY classes. Second, the results depend in an interesting manner on the lead used. For leads $f = 4$ and 5, significantly lower volatility for inflation is available, but at the cost of higher output volatility relative to point A. With lead $f = 6$, it is also possible to obtain points that are Pareto superior to point A in the sense that they produce lower RMSDs of both inflation relative to the target and the output gap. Longer leads (that is, higher f values) tend to produce inferior results.

Figure 5

Some Results with Price-Level Control

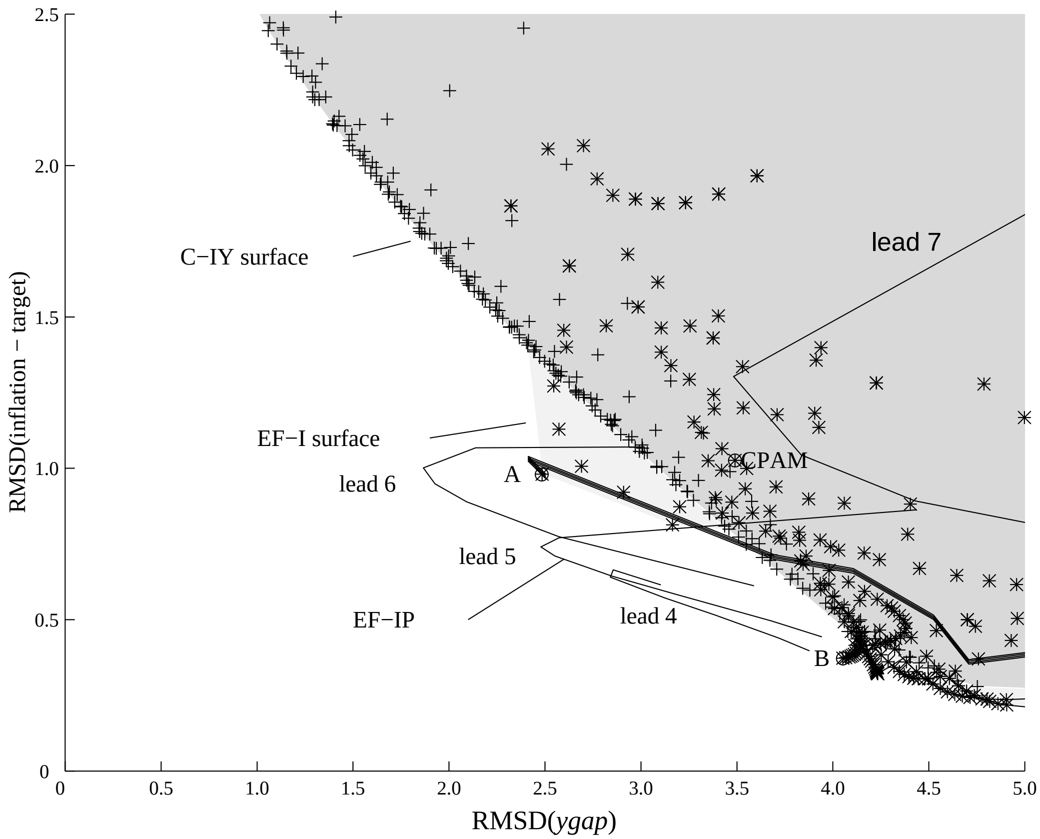

The thick line shown in Figure 5 that starts at point A shows what happens when we add the level element to the reaction function, but with no adjustment to expectations. This line is drawn by varying τ for $f = 6$, and with $\eta = 0$ so there is no adjustment to expectations. This thick line should be compared with the thin line for lead 6 in the figure. The two are markedly different. In particular, while the lead-6 line produces points that are Pareto superior to A, the thick line does not. Starting at point A, the thick line essentially traces out the EF-I frontier, down to the right—lower inflation volatility is available, but at the cost of higher output volatility. In effect, the level controller acts much like the inflation-rate controller in terms of these two measures of performance. Thus, without the adjustment to expectations we get the standard result (as in Fillion and Tetlow 1994 and Haldane and

Figure 6

Frontiers for Policy Rules in the C-IY, EF-I, and EF-IP Classes in Output-Inflation Space

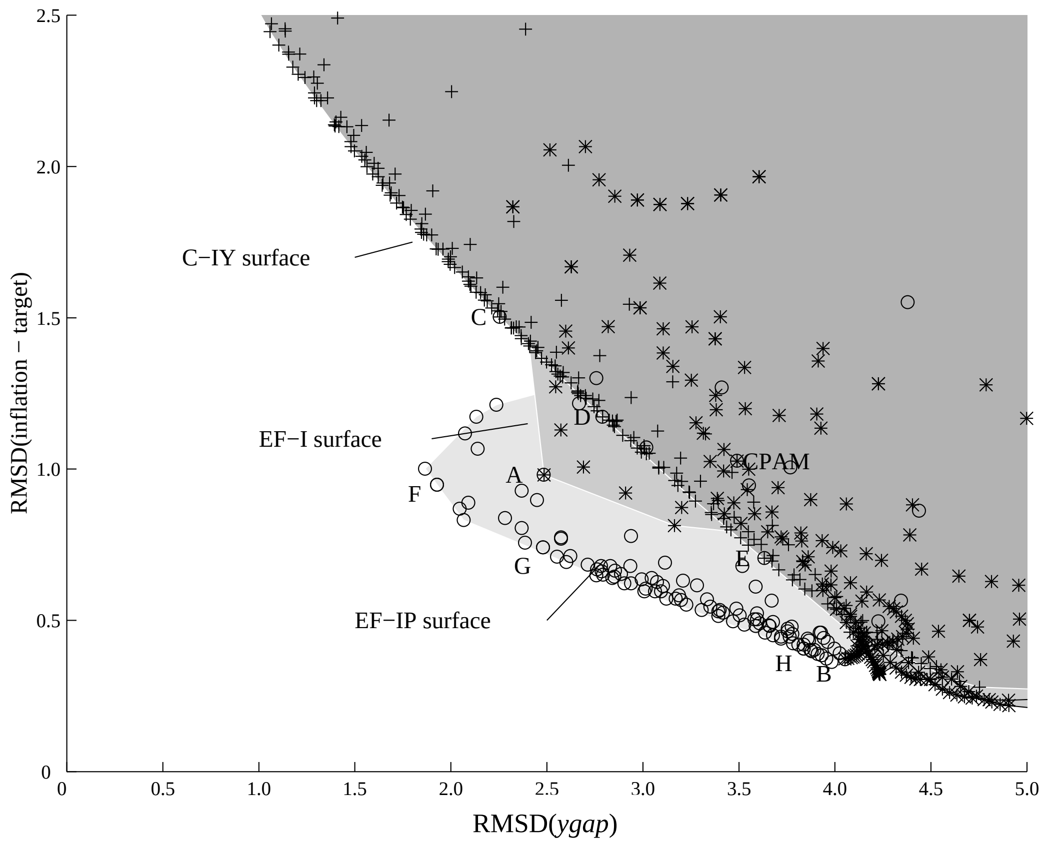

Salmon 1995) that price-level control results in more output variability.[38] This highlights the potentially powerful effects that relatively small adjustments to expectations can have when alternative policy rules are evaluated.

In the second step, we also allow θ, the weight on the inflation-gap terms, to vary, and map out a frontier by varying f and τ. We still hold the

38. Our results with no adjustment to expectations are closer to those of Fillion and Tetlow (1994). Haldane and Salmon (1995) find that a price-level target increases the variability of both inflation and output, while Fillion and Tetlow, like us, find that a price-level target improves inflation control but at the cost of increased output variability. The stronger similarity to Fillion and Tetlow likely reflects the fact that their model and CPAM both specify expectations as a mixture of backward-looking and model-consistent expectations, whereas the Haldane and Salmon model uses only backward-looking expectations.

leads on the inflation terms at their point-A values $(j = 8)$, so the resulting surface is not a globally efficient frontier. But it is good enough for our purposes. As shown in Figure 6, adding the possibility of level control opens up a significant zone of improvement for the EF-IP rule over the EF-I rule. Note, also, that the EF-IP frontier eliminates the local concavity in the EF-I frontier over this region.

In Figure 7, we show the variability of the short-term nominal interest rate versus inflation variability, as in Figure 4. The picture speaks for itself. With a level term in the reaction function, significant reductions in the variability of both nominal interest rates and inflation are available.

Table 6 reports the higher resolution results from longer simulations (128 periods, plus start-up) with 10 replications. The point F gives us the largest reduction in the output gap RMSD relative to point A without raising the RMSD of inflation relative to the target, while point G does the same thing in the other direction. Point H is roughly comparable to point B. We also include the results for the point A in the EF-I class for comparison.

Comparing points A and F we see that the RMSD for the output gap is 0.5 percentage points lower at point F, reflecting both lower output variability and a smaller mean shift at F. In addition, point F has much lower interest rate volatility, as well as better inflation control—the proportion of time outside the +/– 1 percentage point inflation bands falls to 23 per cent as compared with 47 per cent at point A. So, even if there is no preference for reduced price-level variability, point F is the clear winner relative to point A. To the extent that price-level uncertainty is important, point F is even better. At point F, the price-level gap has a mean of 0.2 percentage points with a standard error of 0.5 percentage points, and the average standard deviation of the price-level gap itself is only 1.5 percentage points. Using the point-A rule, these same moments are undefined asymptotically, and for our sample results, they are all an order of magnitude bigger.

Point G provides an interesting alternative to point F. There, the properties of output are about the same as at A, but inflation and interest rate volatility are much lower. It is striking that the proportion of time outside the +/– 1 percentage point bands is now down to 16 per cent. Of course, many points with even better inflation control are available on the surface. Point H is in the region close to point B. At point H, as at point B, exchange rate pass-through is playing a larger role in stabilizing inflation, and inflation control is very precise—inflation is rarely outside the +/– 1 percentage point bands. As with point B, the cost of this (relative to points A, F, and G) is significantly higher output volatility. Note, however, that interest rate volatility is considerably lower at point H than at point B. This illustrates a general conclusion that emerges from this section—that adding some control of the drift in the price level may be an efficient adjunct to a policy of

Figure 7

Frontiers for Policy Rules in the C-IY, EF-I, and EF-IP Classes in Inflation-Interest Rate Space

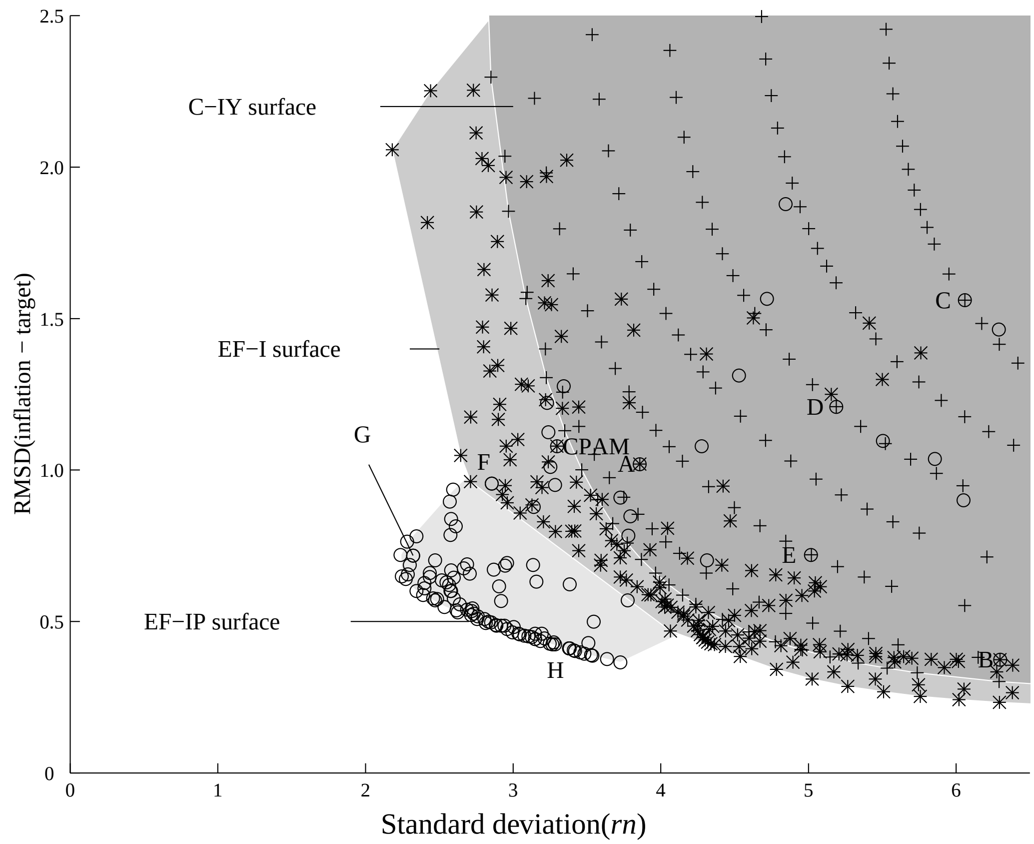

keeping inflation anchored to a low target rate. We stress, however, that this conclusion rests on there being an effect on inflation expectations of the sort we have assumed.

Why do the EF-IP rules work so well? Two mechanisms are at work. The first operates as an automatic stabilizer through real interest rates. Take the case of a negative demand shock that drives prices down. In a world with inflation targeting, the monetary authority will lower short-term nominal interest rates, but to the extent that inflation expectations fall with the observed outcome, the real interest rate effects will be muted; or to put it another way, the monetary authority will have to do more to achieve a given real effect. If the level controller engenders less downward movement in expectations in response to the shock, and eventually an expectation of higher-than-target inflation over some period to re-establish the price-level path, then ex ante real rates will fall (or the monetary authority will have to

Table 6

Comparing Some Alternative Reaction Functions in the EF-IP Family

	Point A	**Point F**	**Point G**	**Point H**
	j = 8 f = 6	j = 8 f = 6	j = 8 f = 5	j = 8 f = 4
	θ = 3.5 τ = 0	θ = 3.5 τ = 2.25	θ = 3.5 τ = 1.75	θ = 1.0 τ = 5.0
RMSD($\pi c_4 - \pi^*$)	1.16	0.85	0.73	0.36
Standard deviation (πc_4)	1.14 (0.19)	0.85 (0.20)	0.73 (0.14)	0.36 (0.03)
Mean($\pi c_4 - \pi^*$)	–0.18 (1.08)	–0.01 (0.05)	–0.01 (0.05)	–0.01 (0.01)
Standard deviation (*pcgap*)	12.21 (7.30)	1.50 (0.47)	1.30 (0.35)	0.36 (0.05)
Mean(*pcgap*)	–2.67 (1.97)	0.23 (0.52)	–0.40 (0.58)	–0.02 (0.06)
RMSD(*ygap*)	2.76	2.25	2.87	4.02
Standard deviation (*ygap*)	2.46 (0.46)	2.19 (0.40)	2.47 (0.54)	3.74 (0.88)
Mean(*ygap*)	–0.77 (1.13)	–0.26 (0.47)	–1.29 (0.99)	–1.37 (0.86)
Standard deviation (*rsl*)	2.56 (0.38)	1.73 (0.37)	1.50 (0.26)	2.45 (0.17)
Standard deviation (*rn*)	4.64 (0.84)	3.15 (0.81)	2.94 (0.71)	4.03 (0.35)
Standard deviation (*z*)	3.51 (0.66)	3.29 (0.60)	3.77 (0.71)	5.02 (0.89)
Per cent of time outside the bands				
+/– 1.0	47	23	16	0
+/– 1.5	30	7	4	0
+/– 2.0	18	2	0	0
Median duration in quarters of departures from bands				
+/– 1.0	10.6	3.3	3.4	0.0
+/– 1.5	9.1	3.0	1.8	0.0
+/– 2.0	6.2	1.5	0.0	0.0

Notes: RMSD is the root-mean-squared deviation. The mean-squared deviation (MSD) for inflation is the variance of inflation plus the squared deviation of inflation from target. The RMSD is the square root of the MSD. As for the other moments, the reported RMSD is calculated by computing the RMSD for each replication of 128 periods, and then averaging over the 10 replications. Standard errors are reported in parentheses.

do less to achieve a given real outcome). This explains why nominal interest rates are less volatile with price-level control.

The second mechanism at work is that inflation depends on expected inflation as well as the output gap, so less variability in inflation and output can be obtained if expectations are more firmly anchored. The mixed backward- and forward-looking model of expectations provides a good way to reconcile forward-looking behaviour with the stylized facts of historical data. Built into this kind of model of expectations are inertia effects in the lags that prolong the process of bringing expectations back to the target, once inflation has moved away from that target. Our extra level term in the expectations equations acts against these inertia effects, providing an outcome where inflation expectations are less extrapolative and revert more quickly to the target. This makes it easier to stabilize inflation, requiring smaller changes in real (as well as nominal) interest rates and hence smaller cycles in output.

Conclusions

In this paper, we investigate the consequences of different forms of policy rules in terms of the trade-offs that emerge for policymakers for the stochastic properties of the economy. The analysis is conducted using a model of the Canadian economy that was constructed for this exercise, which we call CPAM (Canadian Policy Analysis Model). The model has many features patterned on the Bank of Canada's Quarterly Projection Model, QPM, but it has been designed to simulate much faster than QPM. One feature of the model is an asymmetry in inflation dynamics, whereby excess demand creates more upward pressure on inflation than a similar degree of excess supply creates downward pressure. This asymmetry makes the assessment of the performance of alternative rules a bit more complicated than in a linear model, since second moments now in general affect first moments.

We consider six types of disturbances: shocks to (1) potential output (a pure supply shock), (2) commodity prices, (3) the sum of consumption and investment (our domestic demand shock), (4) consumer prices, (5) the exchange rate, and (6) a monetary variable (the slope of the yield curve). The joint distribution of these six shocks is estimated using a VAR for these six variables. Shocks to the structural model are backed out of the VAR by computing the model shocks that reproduce the VAR impulse responses for the first five quarters following a shock. This gives us an implicit covariance structure for model shocks that respects the data and, at the same time, allows us to mimic the very short-run VAR dynamics within the structural model.

The first exercise is to consider the stochastic behaviour of the economy using the base-case CPAM reaction function. This rule is of particular interest because it has been patterned after the reaction function in QPM, which is currently used by Bank staff to provide policy advice. The rule is in the EF-I class and calls for the monetary authority to adjust its instrument on the basis of the expected deviation of inflation from its target about two years in the future. Our simulation results suggest that this rule will produce inflation outcomes that are considerably better than historical outcomes, though roughly in line with the experience in Canada since 1992 when inflation targets came into effect. In particular, using the CPAM rule the model suggests that inflation can be expected to be outside bands of

+/– 1 percentage point around the target about one-third of the time, and the 95 per cent confidence band is about +/– 2 percentage points.

The next exercise considers rules in the EF-I class more generally. Specifically, we compare alternative parameterizations of rules in the EF-I class in terms of their root-mean-squared deviations (RMSDs) of inflation relative to the target and the output gap. The results suggest that, relative to the base-case CPAM rule, there are rules in the EF-I class that produce lower RMSDs of both inflation and output. These rules generally involve a more vigorous reaction to an expected deviation of inflation from the target, and therefore imply more interest rate volatility than the CPAM rule. The results also highlight the fact that in an open economy, there are two channels through which the monetary authority can control inflation—the output gap and the exchange rate—and these different channels produce efficient rules in the EF-I class with quite different properties. Because the lags in the exchange rate channel are shorter than those through the output-gap channel, the monetary authority can achieve substantially better inflation control using the exchange rate channel. This comes at the cost, however, of considerably more variability in interest rates, the exchange rate, and output. Owing to the non-linearity in the Phillips curve in our model, the higher variance of output also results in a lower level of output. Controlling inflation through the output gap has the advantage that stabilizing inflation requires leaning against the output gap, which both reduces output variability and raises the average level of output.

In our simulations with the C-IY class of rules, we find (as have others) that efficient rules in this class have coefficients that are much higher than those suggested by Taylor’s estimated rule characterizing historical U.S. Federal Reserve policy. We also find that, relative to rules in the EF-I class, C-IY rules produce roughly similar results if we limit the comparison to the RMSDs of inflation and output. When we look deeper, however, the C-IY rules seem less appealing. First, this class of rules is not generally capable of achieving the inflation target on average in CPAM without an explicit modification to allow for the shift in the mean of output discussed above. Our results have penalized this by using the square of the bias in the RMSD, but this is not the same as ensuring the target is met on average. Second, the volatility of interest rates is found to be relatively high with C-IY rules, sometimes high enough in absolute terms that a considerable proportion of observations have negative nominal interest rates. With EF rules, where the dynamic structure of the model is taken into account, efficient results are achieved with instrument settings that are much less volatile.

Finally, we turn to the question of whether there is any possible merit in including a reaction to the drift in the price level in the monetary rule. The conventional view is that, all else being equal, output variability will rise if a level target is adopted. This result also emerges in our model. However, when we add an element of price-level control while taking the implications of the Lucas critique seriously, the results change dramatically. In particular, when we allow expectations to adjust to take account of the effects of price-level control, we find that introducing an element of price-level control results in a significant improvement in the trade-offs facing the policymaker. A significant new region is identified with the EF-IP class of rules that dominates the EF-I and C-IY classes for a wide range of preferences. It is possible to achieve reductions in volatility in all dimensions, and to raise the mean level of output as a result. We stress that these results are only suggestive. However, we think that they establish a case for looking seriously at the possible role of level conditions in the design of efficient policy rules for a world of price stability.

As in any project of this scope, we have been required to make myriad specification decisions about the model, about the shocks, and about the reaction functions we use. In future work, the robustness of our results to plausible variations in the structure of the model, to alternative characterizations of the shocks, and to larger replications of our experiments should all be considered. A number of variations on the reaction functions we study also merit consideration. Although we find EF-I rules generally outperform C-IY rules, explicitly forward-looking rules that put some weight on output as well as inflation performance (such as, EF-IY) deserve further attention. The important effects of exchange rate pass-through on consumer prices in an economy such as Canada's also suggest that it would be interesting to consider rules based on an aggregate price index that excludes (or reduces) the direct effect of exchange rate changes, such as the model's numeraire price, the domestic absorption deflator.

Finally, it must be emphasized that our stochastic simulation analysis provides only the simplest conceivable extension into the realm of uncertainty and monetary policy. We take shocks seriously. But we maintain the heroic assumption that the monetary authority knows the true data-generating process. So do private agents, in the main, though we do not use fully model-consistent expectations for inflation. Of great importance, also, is the assumption that private agents know the monetary rule. This eliminates any problems of uncertainty regarding the monetary regime and the possible confusion of the effects of shocks with perceptions of regime change. All these issues merit consideration in future research.

Appendix 1

Defining the Shocks to the Structural Model

This appendix describes precisely how we define the shocks to the structural model from the results of the estimated VAR, such that the model can produce impulse response functions that are, in a sense described below, close to those from the VAR. This done, the model can be stochastically simulated by subjecting it to random collections of these impulses.

To understand how we do this, it is convenient to start with a simple example, and then to add various complications, such as "forward-lookingness" and non-linearities. Suppose that the VAR is written as

$$x_t = A(L)x_{t\text{-}1} + \Omega\varepsilon_t, \qquad \varepsilon{\sim}N(0, I), \tag{A1}$$

where x is the column vector $[y^*, pcom, c + i, py, z, rsl]$, A denotes the matrix of coefficients, L is the lag operator, $\Omega^T\Omega$ is the variance-covariance matrix of the VAR, ε is the vector of iid shocks, and I denotes an identity matrix. In order to generate the jth impulse response functions in Figure 1, the iid shock term in equation (A1) is set to ι^j, where ι^j is a vector of zeros except in the jth row where it is 1. Because the VAR is linear in its shock terms, the stochastic driving force may be viewed as coming from the six impulse response functions rather than from the single multivariate shock term ε. That is,

$$x_t = A(L)x_{t\text{-}1} + \Omega\iota^1\varepsilon_{1t} + \ldots + \Omega\iota^6\varepsilon_{6t}, \qquad \varepsilon{\sim}N(0, I). \tag{A2}$$

Note that each term $\Omega\iota^j\varepsilon_{jt}$ corresponds to the shock performed to produce one of the impulse response functions, weighted by an $N(0, 1)$ random number. This means that, in order to generate equation (A1) stochastically, it is sufficient to generate six single-period random shocks and then apply a single aggregate shock equal to the sum of these six impulse responses multiplied by the random shock.

To use this idea in the context of CPAM, a forward-looking non-linear model, the first step is to work out how to generate the impulse functions. The second step is to construct single-period shock terms to generate these impulses and to attach random numbers to these shock terms, as in equation (A2).

The first of these steps, how to generate the VAR impulse responses in the model, is relatively straightforward. In CPAM, each behavioural equation is associated with a shock term, and the model can be simulated with a shock term attached to each of the behavioural equations for $\{y^*, pcom, c+i, py, z, rsl\}$ and the actual variables $\{y^*, pcom, c+i, py, z, rsl\}$ set exogenously with their value set equal to the impulse response for each variable. That is, the shock terms necessary to replicate each impulse response function can be "backed out" of the model. By performing six such simulations, one for each impulse response function from the VAR, this method, by construction, gives the shock terms necessary to replicate each one of the VAR's impulse response functions. Note that since each one of these simulations produces six shock terms there are actually 36 shock series produced. Later on, transformed values of these shock series will correspond (roughly speaking) to the 36 elements in the matrix Ω.

One minor complication is that the shock series do not generally return to zero. Although in most cases the VAR impulse responses correspond closely with the model's properties, so the shock terms settle down to a small number, they do not go to exactly zero. We therefore truncate the shock series by setting them to zero after five quarters. The five-quarter horizon is motivated by the fact that beyond five quarters the monetary reaction begins to have a significant effect on the economy.[1]

The next stage of the exercise is to take these truncated shock series, which in general have some serial correlations, and rewrite them to be generated by single-period shocks (as in equation A2).

Suppose the series for the ith shock term necessary to replicate the jth impulse response is given by $(\mu^i_{j0}, \mu^i_{j1}, \ldots, \mu^i_{j4}, 0, 0, \ldots)$ (relative to some initial starting date) after it has been truncated as described above. Then this can be generated by the single period shock ε^i_{jt} and the equation

$$\mu^i_{jt} = \mu^i_{j0}\varepsilon^i_{jt} + \mu^i_{j1}\varepsilon^i_{jt-1} + \ldots + \mu^i_{j4}\varepsilon^i_{jt-4}. \tag{A3}$$

At this point, we have a collection of 36 (truncated) shock series. The last step of the exercise is to combine them in such a way that they generate the six model shock terms. This is done by noting that the shock term to the behavioural equation corresponding to the variable i, μ^i_t, is given by

$$\mu^i_t = \mu^i_{1t} + \ldots + \mu^i_{6t}, \tag{A4}$$

the sum of the six shock terms from the VAR corresponding to the ith variable. The complete system can then be written as:

$$\mu^i_t = \sum_{j=1}^{6} \sum_{k=0}^{4} \mu^i_{jk}\varepsilon^i_{jt-k}, \qquad 1 \le i \le 6. \tag{A5}$$

1. The results do not appear overly sensitive to this choice of a truncation horizon.

References

Armstrong, J., R. Black, D. Laxton, and D. E. Rose. 1995. *The Bank of Canada's New Quarterly Projection Model.* Part 2, *A Robust Method for Simulating Forward-Looking Models*. Technical Report No. 73. Ottawa: Bank of Canada.

Ball, L. 1997. "Efficient Rules for Monetary Policy." NBER Working Paper No. 5952. National Bureau of Economic Research, Cambridge, Mass.

Bean, C. 1996. "The Convex Phillips Curve and Macroeconomic Policymaking Under Uncertainty." Unpublished, November. London School of Economics and H.M. Treasury.

Black, R., D. Laxton, D. E. Rose, and R. Tetlow. 1994. *The Bank of Canada's New Quarterly Projection Model.* Part 1, *The Steady-State Model: SSQPM.* Technical Report No. 72. Ottawa: Bank of Canada.

Black, R., and D. Rose. 1997. "The Canadian Policy Analysis Model: CPAM." Working Paper No. 97-16. Bank of Canada, Ottawa.

Black, R., D. Coletti, and S. Monnier. 1998. "On the Costs and Benefits of Price Stability." Bank of Canada. This volume.

Blanchard, O. J. 1985. "Debt, Deficits, and Finite Horizons." *Journal of Political Economy* 93 (April): 223-47.

Bryant, R. C., P. Hooper, and C. L. Mann, editors. 1993. *Evaluating Policy Regimes: New Research in Empirical Macroeconomics.* Washington, D.C.: Brookings Institution.

Buiter, W. H., and M. H. Miller. 1985. "Costs and Benefits of an Anti-Inflationary Policy: Questions and Issues." In *Inflation and Unemployment: Theory, Experience and Policy-Making*, edited by V. E. Argy and J. W. Nevile, 11-38. London: Allen & Unwin.

Clark, P. B., and D. Laxton. 1997. "Phillips Curves, Phillips Lines and the Unemployment Costs of Overheating." IMF Working Paper No. 97/17. International Monetary Fund, Washington, D.C.

Clark, P., D. Laxton, and D. Rose. 1996. "Asymmetry in the U.S. Output-Inflation Nexus." *International Monetary Fund Staff Papers* 43 (March): 216-51.

———. 1997. "Capacity Constraints, Inflation and the Transmission Mechanism: Forward-Looking Versus Myopic Policy Rules." Unpublished, revision of IMF Working Paper No. 95/75. International Monetary Fund, Washington, D.C.

Cogley, T., and J. M. Nason. 1995. "The Effects of the Hodrick-Prescott Filter on Trend and Difference Stationary Time Series: Implications for Business Cycle Research." *Journal of Economic Dynamics and Control* 19 (January–February): 253-78.

Coletti, D., B. Hunt, D. Rose, and R. Tetlow. 1996. *The Bank of Canada's New Quarterly Projection Model.* Part 3, *The Dynamic Model: QPM*. Technical Report No. 75. Ottawa: Bank of Canada.

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

Côté, A., and T. Macklem. 1996. "The Determination of Interest Rates and the Exchange Rate in the Bank of Canada's Quarterly Projection Model." In *The Determination of Long-Term Interest Rates and Exchange Rates and the Role of Expectations.* Conference Papers Vol. 2, 252-73. Basle: Bank for International Settlements.

Crawford, A., and M. Kasumovich. 1996. "Does Inflation Uncertainty Vary with the Level of Inflation?" Working Paper No. 96-9. Bank of Canada, Ottawa.

Debelle, G., and D. Laxton. 1996. "Is the Phillips Curve Really a Curve? Some Evidence for Canada, the United Kingdom, and the United States." IMF Working Paper No. 96/111. International Monetary Fund, Washington, D.C.

Duguay, P. 1994. "Some Thoughts on Price Stability Versus Zero Inflation." Paper presented at the conference on Central Bank Independence and Accountability, Università Bocconi, Milan, March 1994.

Dupasquier, C., and N. Ricketts. 1998. "Non-Linearities in the Output-Inflation Relationship." Bank of Canada. This volume.

Fillion, J.-F., and R. Tetlow. 1994. "Zero-Inflation or Price-Level Targeting? Some Answers from Stochastic Simulations on a Small Open-Economy Macro Model." in *Economic Behaviour and Policy Choice Under Price Stability*, 129-66. Proceedings of a conference held at the Bank of Canada, October 1993. Ottawa: Bank of Canada.

Fisher, P. G., L. Mahadeva, and J. D. Whitley. 1996. "The Output Gap and Inflation: Experience at the Bank of England." Paper prepared for the BIS Model Builders' Meeting, Basle (January).

Frenkel, J. A., and A. Razin. 1992. *Fiscal Policies and the World Economy.* 2nd ed. Cambridge, Mass.: MIT Press.

Fuhrer, J. C. 1994. "Optimal Monetary Policy in a Model of Overlapping Price Contracts." Working Paper No. 94-2. Federal Reserve Bank of Boston.

Fuhrer, J. C., and G. R. Moore. 1995a. "Inflation Persistence." *Quarterly Journal of Economics* 110 (February): 127-59.

———. 1995b. "Monetary Policy Trade-offs and the Correlation between Nominal Interest Rates and Real Output." *American Economic Review* 85 (March): 219-39.

Grimes, A., and J. Wong. 1994. "The Role of the Exchange Rate in New Zealand Monetary Policy." In *Exchange Rate Policy and Interdependence: Perspectives from the Pacific Basin,* edited by R. Glick and M. Hutchison, 181-95. New York: Cambridge University Press.

Guay, A., and P. St-Amant. 1996. *Do Mechanical Filters Provide a Good Approximation of Business Cycles?* Technical Report No. 78. Ottawa: Bank of Canada.

Haldane, A. G., and C. K. Salmon. 1995. "Three Issues on Inflation Targets." In *Targeting Inflation*, edited by A. Haldane, 170-201. London: Bank of England.

Henderson, D. W., and W. J. McKibbin. 1993. "A Comparison of Some Basic Monetary Policy Regimes for Open Economies: Implications of Different Degrees of Instrument Adjustment and Wage Persistence." *Carnegie-Rochester Conference Series on Public Policy* 39 (December): 221-317.

Isard, P., and D. Laxton. 1996. "Strategic Choice in Phillips Curve Specification: What if Bob Gordon is Wrong?" Paper presented at the European University Institute Conference on European Unemployment: Macroeconomic Aspects, Florence, Italy.

Judd, J. P., and B. Motley. 1991. "Nominal Feedback Rules for Monetary Policy." *Federal Reserve Bank of San Fransisco Economic Review* (Summer): 3-17.

Laxton, D., G. Meredith, and D. Rose. 1995. "Asymmetric Effects of Economic Activity on Inflation: Evidence and Policy Implications." *International Monetary Fund Staff Papers* 42 (June): 344-74.

Laxton, D., D. Rose, and R. Tetlow. 1993a. "Is the Canadian Phillips Curve Non-Linear?" Working Paper No. 93-7. Bank of Canada, Ottawa.

———. 1993b. *Monetary Policy, Uncertainty and the Presumption of Linearity.* Technical Report No. 63. Ottawa: Bank of Canada.

Lebow, D. E., J. M. Roberts, and D. J. Stockton. 1992. "Economic Performance Under Price Stability," Working Paper No. 125, Division of Research and Statistics, U.S. Board of Governors of the Federal Reserve System, Washington, D.C.

Levin, A. 1996. "A Comparison of Alternative Monetary Policy Rules in the Federal Reserve Board's Multi-Country Model." In *The Determination of Long-Term Interest Rates and Exchange Rates and the Role of Expectations,* 340-66. Basle: Bank for International Settlements.

Macklem, T., A. Paquet, and L. Phaneuf. 1996. "Asymmetric Effects of Monetary Policy: Evidence from the Yield Curve." Working Paper No. 42. Center for Research on Economic Fluctuations and Employment, Université de Québec à Montréal.

McCallum, B. 1984. "Monetarist Rules in the Light of Recent Experience." *American Economic Review* 74 (May): 388-91.

———. 1997. *Issues in the Design of Monetary Policy Rules*. Cambridge, Mass.: National Bureau of Economic Research.

Meade, J. E. 1978. "The Meaning of 'Internal Balance.' " *Economic Journal* 88 (September): 423-35.

Phillips, A. W. 1958. "The Relation Between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861-1957." *Economica* 25 (November): 283-99.

Sargent, T. J. 1971. "A Note on the 'Accelerationist' Controversy." *Journal of Money, Credit, and Banking* 3 (August): 721-25.

Stevens, G., and G. Debelle. 1995. "Monetary Policy Goals for Inflation in Australia." In *Targeting Inflation*, edited by A. Haldane, 81-100. London: Bank of England.

Summers, L. 1991. "How Should Long-Term Monetary Policy Be Determined? Panel Discussion." *Journal of Money, Credit, and Banking* 23 (August): 625-31.

Svensson, L. 1996a. "Price Level Targeting vs. Inflation Targeting: A Free Lunch?" NBER Working Paper No. 5719. National Bureau of Economic Research, Cambridge, Mass.

———. 1996b. "Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets." Working Paper Series No. 56. Bank of England, London.

Taylor, J. B. 1979. "Estimation and Control of a Macroeconomic Model with Rational Expectations." *Econometrica* 47 (September): 1267-86.

———. 1980. "Aggregate Dynamics and Staggered Contracts." *Journal of Political Economy* 88 (February): 1-23.

———. 1993. "Discretion Versus Policy Rules in Practice." *Carnegie-Rochester Conference Series on Public Policy* 39 (December): 195-214.

———. 1994. "The Inflation-Output Variability Tradeoff Revisited." In *Goals, Guidelines, and Constraints Facing Monetary Policymakers*, edited by J. C. Fuhrer, 21-38. Conference Series No. 38. Boston: Federal Reserve Bank of Boston.

Tetlow, R., and P. von zur Muehlen. 1996. "Monetary Policy Rules in a Small Forward-Looking Macro Model." Unpublished. U.S. Board of Governors of the Federal Reserve System, Washington, D.C.

Tobin, J. 1980. "Stabilization Policy Ten Years After." *Brookings Papers on Economic Activity* 1: 19-71.

Turner, D. 1995. "Speed Limit and Asymmetric Inflation Effects from the Output Gap in the Major Seven Economies." *OECD Economic Studies* 24 (1995/I): 57-87.

———. 1996. "Inflation Targeting in New Zealand: Is a Two Per Cent Band Width Feasible?" Unpublished, March. Economics Department, OECD.

Weil, P. 1989. "Overlapping Families of Infinitely-Lived Agents." *Journal of Public Economics* 38 (March): 183-98.