Volume Title: Monetary Policy Rules

Volume Author/Editor: John B. Taylor, editor

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-79124-6

Volume URL: http://www.nber.org/books/tayl99-1

Publication Date: January 1999

Chapter Title: Policy Rules for Open Economies

Chapter Author: Laurence M. Ball

Chapter URL: http://www.nber.org/chapters/c7415

Chapter pages in book: (p. 127 - 156)

3 Policy Rules for Open Economies

Laurence Ball

3.1 Introduction

What policy rules should central banks follow? A growing number of economists and policymakers advocate targets for the level of inflation. Many also argue that inflation targeting should be implemented through a "Taylor rule" in which interest rates are adjusted in response to output and inflation. These views are supported by the theoretical models of Svensson (1997) and Ball (1997), in which the optimal policies are versions of inflation targets and Taylor rules.

Many analyses of policy rules assume a closed economy. This paper extends the Svensson-Ball model to an open economy and asks how the optimal policies change. The short answer is they change quite a bit. In open economies, inflation targets and Taylor rules are suboptimal unless they are modified in important ways. Different rules are required because monetary policy affects the economy through exchange rate as well as interest rate channels.[1]

Section 3.2 presents the model, which consists of three equations. The first is a dynamic, open economy IS equation: output depends on lags of itself, the real interest rate, and the real exchange rate. The second is an open economy Phillips curve: the change in inflation depends on lagged output and the lagged change in the exchange rate, which affects inflation through import prices. The

Laurence Ball is professor of economics at Johns Hopkins University and a research associate of the National Bureau of Economic Research.

The author is grateful for research assistance from Qiming Chen, Nada Choueiri, and Heiko Ebens, and for comments and advice from Bob Buckle, Pedro deLima, Eric Hansen, David Gruen, David Longworth, Tiff Macklem, Scott Roger, Thomas Sargent, John Taylor, and seminar participants at the Federal Reserve Bank of Kansas City and the Reserve Banks of Australia and Canada.

1. Svensson (forthcoming) also examines alternative policy rules in an open economy model. That paper differs from this one and from Svensson (1997) in stressing microfoundations and forward-looking behavior, at the cost of greater complexity.

final equation is a relation between interest rates and exchange rates that captures the behavior of asset markets.

Section 3.3 derives the optimal instrument rule in the model. This rule differs in two ways from the Taylor rule that is optimal in a closed economy. First, the policy instrument is a weighted sum of the interest rate and the exchange rate—a "monetary conditions index" like the ones used in several countries. Second, on the right-hand side of the policy rule, inflation is replaced by "long-run" inflation. This variable is a measure of inflation adjusted for the temporary effects of exchange rate fluctuations.

Section 3.4 considers several instrument rules proposed in other papers at this conference. I find that most of these rules perform poorly in my model.

Section 3.5 turns to inflation targeting. In the closed economy models of Svensson and Ball, a simple version of this policy is equivalent to the optimal Taylor rule. In an open economy, however, inflation targeting can be dangerous. The reason concerns the effects of exchange rates on inflation through import prices. This is the fastest channel from monetary policy to inflation, and so inflation targeting implies that it is used aggressively. Large shifts in the exchange rate produce large fluctuations in output.

Section 3.6 presents a more positive result. While pure inflation targeting has undesirable effects, a modification produces much better outcomes. The modification is to target "long-run" inflation—the inflation variable that appears in the optimal instrument rule. This variable is not influenced by the exchange-rate-to-import-price channel, and so targeting it does not induce large exchange rate movements. Targeting long-run inflation is not exactly equivalent to the optimal instrument rule, but it is a close approximation for plausible parameter values.

Section 3.7 concludes the paper.

3.2 The Model

3.2.1 Assumptions

The model is an extension of Svensson (1997) and Ball (1997) to an open economy. The goal is to capture conventional wisdom about the major effects of monetary policy in a simple way. The model is similar in spirit to the more complicated macroeconometric models of many central banks.

The model consists of three equations:

$$y = -\beta r_{-1} - \delta e_{-1} + \lambda y_{-1} + \varepsilon, \tag{1}$$

$$\pi = \pi_{-1} + \alpha y_{-1} - \gamma(e_{-1} - e_{-2}) + \eta, \tag{2}$$

$$e = \theta r + v, \tag{3}$$

where y is the log of real output, r is the real interest rate, e is the log of the real exchange rate (a higher e means appreciation), π is inflation, and ε, η, and

v are white noise shocks. All parameters are positive, and all variables are measured as deviations from average levels.

Equation (1) is an open economy IS curve. Output depends on lags of the real interest rate and the real exchange rate, its own lag, and a demand shock.

Equation (2) is an open economy Phillips curve. The change in inflation depends on the lag of output, the lagged change in the exchange rate, and a shock. The change in the exchange rate affects inflation because it is passed directly into import prices. This interpretation is formalized in the appendix, which derives equation (2) from separate equations for domestic goods and import inflation.

Finally, equation (3) posits a link between the interest rate and the exchange rate. It captures the idea that a rise in the interest rate makes domestic assets more attractive, leading to an appreciation. The shock v captures other influences on the exchange rate, such as expectations, investor confidence, and foreign interest rates. Equation (3) is similar to reduced-form equations for the exchange rate in many textbooks.

The central bank chooses the real interest rate r. One can interpret any policy rule as a rule for setting r. Using equation (3), one can also rewrite any rule as a rule for setting e, or for setting some combination of e and r.[2]

A key feature of the model is that policy affects inflation through two channels. A monetary contraction reduces output and thus inflation through the Phillips curve, and it also causes an appreciation that reduces inflation directly. The lags in equations (1), (2), and (3) imply that the first channel takes two periods to work: a tightening raises r and e contemporaneously, but it takes a period for these variables to affect output and another period for output to affect inflation. In contrast, the direct effect of an exchange rate change on inflation takes only one period. These assumptions capture the common view that the direct exchange rate effect is the quickest channel from policy to inflation.

3.2.2 Calibration

In analyzing the model, I will interpret a period as a year. With this interpretation, the time lags in the model are roughly realistic. Empirical evidence suggests that policy affects inflation through the direct exchange rate channel in about a year, and through the output channel in about two years (e.g., Reserve Bank of New Zealand 1996; Lafleche 1996).

The analysis will use a set of base parameter values. Several of these values are borrowed from the closed economy model in Ball (1997). Based on evidence discussed there, I assume that λ, the output persistence coefficient, is 0.8; that α, the slope of the Phillips curve, is 0.4; and that the total output loss from a 1-point rise in the interest rate is 1.0. In the current model, this total

2. Note I assume that policymakers set the *real* interest rate. In practice, the interest rates controlled directly by policymakers are nominal rates. However, policymakers can move the real interest rate to their desired level by setting the nominal rate equal to the desired real rate plus inflation.

effect is $\beta + \delta\theta$: β is the direct effect of the interest rate and $\delta\theta$ is the effect through the exchange rate. I therefore assume $\beta + \delta\theta = 1.0$.

The other parameters depend on the economy's degree of openness. My base values are meant to apply to medium to small open economies such as Canada, Australia, and New Zealand. My main sources for the parameters are studies by these countries' central banks. I assume $\gamma = 0.2$ (a 1 percent appreciation reduces inflation by two-tenths of a point) and $\theta = 2.0$ (a 1 point rise in the interest rate causes a 2 percent appreciation). I also assume $\beta/\delta = 3.0$, capturing a common rule of thumb about IS coefficients. Along with my other assumptions, this implies $\beta = 0.6$ and $\delta = 0.2$.[3]

3.3 Efficient Instrument Rules

Following Taylor (1994), the optimal policy rule is defined as the one that minimizes a weighted sum of output variance and inflation variance. The weights are determined by policymakers' tastes. As in Ball (1997), an "efficient" rule is one that is optimal for some weights, or equivalently a rule that puts the economy on the output-inflation variance frontier. This section derives the set of efficient rules in the model.

3.3.1 Variables in the Rule

As discussed earlier, we can interpret any policy rule as a rule for r, a rule for e, or a rule for a combination of the two. Initially, it is convenient to consider rules for e. To derive the efficient rules, I first substitute equation (3) into equation (1) to eliminate r from the model. I shift the time subscripts forward to show the effects of the current exchange rate on future output and inflation. This yields

$$y_{+1} = -(\beta/\theta + \delta)e + \lambda y + \varepsilon_{+1} + (\beta/\theta)v, \tag{4}$$

$$\pi_{+1} = \pi + \alpha y - \gamma(e - e_{-1}) + \eta_{+1}. \tag{5}$$

Consider a policymaker choosing the current e. One can define the state variables of the model by two expressions corresponding to terms on the right-hand sides of equations (4) and (5): $\lambda y + (\beta/\theta)v$ and $\pi + \alpha y + \gamma e_{-1}$. The future paths of output and inflation are determined by these two expressions, the rule for choosing e, and future shocks. Since the model is linear-quadratic, one can show the optimal rule is linear in the two state variables:

$$e = m[\lambda y + (\beta/\theta)v] + n(\pi + \alpha y + \gamma e_{-1}), \tag{6}$$

where m and n are constants to be determined.

In equation (6), the choice of e depends on the exchange rate shock v as

3. Examples of my sources for base parameter values are the Canadian studies of Longworth and Poloz (1986) and Duguay (1994) and the Australian study of Gruen and Shuetrim (1994).

well as observable variables. By equation (3), v can be replaced by $e - \theta r$. Making this substitution and rearranging terms yields

$$(7) \qquad wr + (1 - w)e = ay + b(\pi + \gamma e_{-1}),$$

where

$$w = m\beta\theta/(\theta - m\beta + m\beta\theta), \qquad a = \theta(m\lambda + n\alpha)/(\theta - m\beta + m\beta\theta),$$

$$b = n\theta/(\theta - m\beta + m\beta\theta).$$

This expresses the optimal policy as a rule for an average of r and e.

3.3.2 Interpretation

In the closed economy model of Svensson and Ball, the optimal policy is a Taylor rule: the interest rate depends on output and inflation. Equation (7) modifies the Taylor rule in two ways. First, the policy variable is a combination of r and e. And second, inflation is replaced by $\pi + \gamma e_{-1}$, a combination of inflation and the lagged exchange rate. Each of these modifications has a simple interpretation.

The first result supports the practice of using an average of r and e—a "monetary conditions index" (MCI)—as the policy instrument. Several countries follow this approach, including Canada, New Zealand, and Sweden (see Gerlach and Smets 1996). The rationale for using an MCI is that it measures the overall stance of policy, including the stimulus through both r and e. Policymakers shift the MCI when they want to ease or tighten. When there are shifts in the e/r relation—shocks to equation (3)—r is adjusted to keep the MCI at the desired level.

The second modification of the Taylor rule is more novel. The term $\pi + \gamma e_{-1}$ can be interpreted as a long-run forecast of inflation under the assumption that output is kept at its natural level. With a closed economy Phillips curve, this forecast would simply be current inflation. In an open economy, however, inflation will change because the exchange rate will eventually return to its long-run level, which is normalized to zero. For example, if e was positive in the previous period, there will be a depreciation of e_{-1} at some point starting in the current period. By equation (2), this will raise inflation by γe_{-1} at some point after the current period. I will use the term "long-run inflation" and the symbol π^* to stand for $\pi + \gamma e_{-1}$.

More broadly, one can interpret $\pi + \gamma e_{-1}$ as a measure of inflation that filters out direct but temporary effects of the exchange rate. For a given output path, an appreciation causes inflation to fall, but it will rise again by γe_{-1} when the appreciation is reversed. The adjustment from π to π^* is similar in spirit to calculations of "core" or "underlying" inflation by central banks. These variables are measures of inflation adjusted for transitory influences such as changes in indirect taxes or commodity prices. Many economists argue that policy should respond to underlying inflation and ignore transitory fluctuations. My model supports this idea for the case of fluctuations caused by exchange rates.

3.3.3 Efficient Coefficients for the Rule

The coefficients in the policy rule (7) depend on the constants m and n, which are not yet determined. The next step is to derive the efficient combinations of m and n—the combinations that put the economy on the output-inflation variance frontier. As discussed in the appendix, the set of efficient policies depends on the coefficients in equations (1), (2), and (3) but not on the variances of the three shocks (although these variances determine the absolute position of the frontier). For base parameter values, I compute the variances of output and inflation for given m and n and then search for combinations that define the frontier.

Figure 3.1 presents the results in a graph. The figure shows the output-inflation variance frontier when the variance of each shock is one. For selected points on the frontier, the graph shows the policy rule coefficients that put the economy at that point. It also shows the weights on output variance and inflation variance that make each policy optimal.

Two results are noteworthy. The first concerns the weights on r and e in the MCI. There is currently a debate among economists about the appropriate weights. Some argue that the weights should be proportional to the coefficients on e and r in the IS equation (e.g., Freedman 1994). For my base parameters, this implies $w = 0.75$, that is, weights of 0.75 on r and 0.25 on e. Others suggest a larger weight on e to reflect the direct effect of the exchange rate on inflation (see Gerlach and Smets 1996). In my model, the optimal weight on e is larger than 0.25, but by a small amount. For example, if the policymaker's objective function has equal weights on output and inflation variances, the MCI weight on e is 0.30. The weight on e is much smaller than its relative short-run effect on inflation. The only exceptions occur when policymakers' objectives have very little weight on output variance.[4]

The second result concerns the coefficients on y and π^*, and how they compare to the optimal coefficients on y and π in a closed economy. Note that a 1 point rise in the interest rate, which also raises the exchange rate, raises the MCI by a total of $w + \theta(1 - w)$. Dividing the coefficients on y and π^* by this expression yields the responses of r to movements in y and π^* (holding constant the exchange rate shock v). These responses are the analogues of Taylor rule coefficients in a closed economy. For equal weights in policymakers' objective functions, the interest rate response to output is 1.04 and the response to π^* is 0.82. Assuming the same objective function, the corresponding re-

4. One measure of the overall effect of e on inflation is the effect through appreciation in one period plus the effect through the Phillips curve in two periods. This sum is $\gamma + \delta\alpha = 0.28$. The corresponding effect of r on inflation is $\beta\alpha = 0.24$. The MCI would put more weight on e than on r if it were based on these inflation effects.

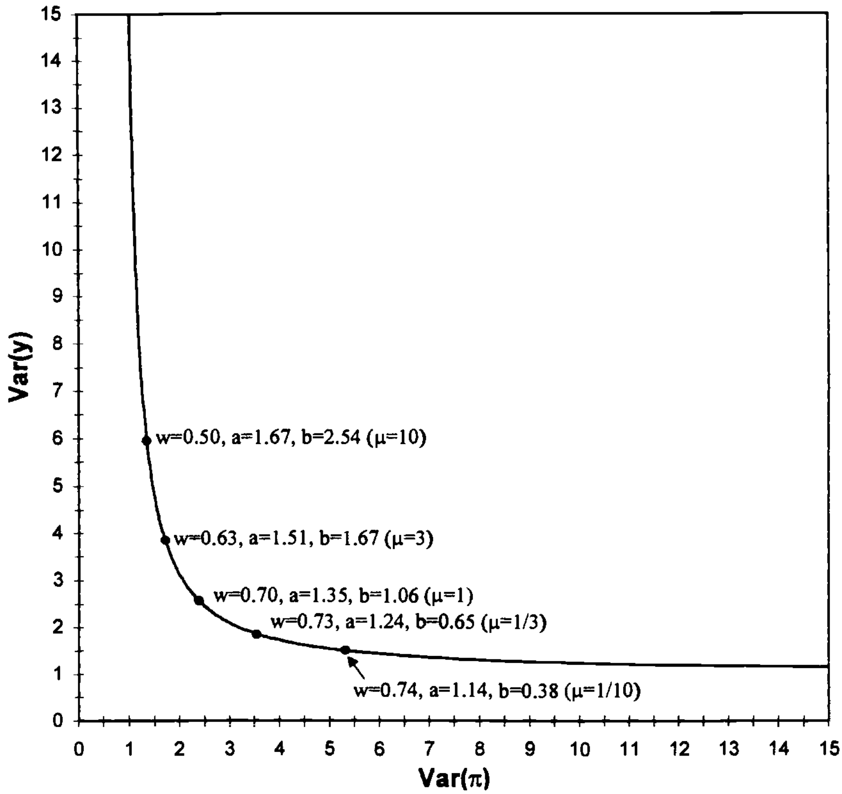

Fig. 3.1 Output-inflation variance frontier
Note: Objective function = var(y) + μ var(π).

sponses in a closed economy are 1.13 for output and 0.82 for inflation (Ball 1997). Thus the sizes of interest rate movements are similar in the two cases.

3.4 Other Instrument Rules

This paper is part of a project to evaluate policy rules in alternative macroeconomic models. As part of the project, all authors are evaluating a list of six rules to see whether any performs well across a variety of models. Each of the rules has the general form

$$r = ay + b\pi + cr_{-1}, \tag{8}$$

where a, b, and c are constants. Table 3.1 gives the values of the constants in the six rules.

All of these rules are inefficient in the current model. There are two separate

Table 3.1 **Alternative Policy Rules**

	Rule					
	1	2	3	4	5	6
a	0.8	1	0.5	1	0.06	0.08
b	2	0.2	0.5	0.5	0.2	0.3
c	1	1	0	0	2.86	2.86
Base case						
var(y)	531.59	4.42	2.62	1.86	∞	∞
var(π)	5.18	6.55	3.43	4.05	∞	∞
Closed economy						
var(y)	∞	6.53	2.77	1.81	∞	∞
var(π)	∞	7.59	3.91	4.22	∞	∞

problems. First, the rules are designed for closed economies and therefore do not make the adjustments for exchange rate effects discussed in the last section. Second, even if the economy were closed, the coefficients in most of the rules would be inefficient. To distinguish between these problems, I evaluate the rules in two versions of my model: the open economy case considered above and a closed economy case obtained by setting δ and γ to zero. The latter is identical to the model in Ball (1997).[5]

Table 3.1 presents the variances of output and inflation for the six rules. The rules fall into two categories. The first are those with c, the coefficient on lagged r, equal to or greater than one (rules 1, 2, 5, and 6). For these rules, the output and inflation variances range from large to infinite, both in closed and open economies. This result reflects the fact that efficient rules in either case do not include the lagged interest rate. Including this variable leads to inefficient oscillations in output and inflation.

The other rules, numbers 3 and 4, omit the lagged interest rate ($c = 0$). These rules perform well in a closed economy. Indeed, rule 4 is fully efficient in that case; rule 3 is not quite efficient, but it puts the economy close to the frontier (see Ball 1997). In an open economy, however, rules 3 and 4 are inefficient because they ignore the exchange rate. Rule 4, for example, produces an output variance of 1.86 and an inflation variance of 4.05. Using an efficient rule, policy can achieve the same output variance with an inflation variance of 3.54.

Recall that the set of efficient rules does not depend on the variances of the model's three shocks. In contrast, the losses from using an inefficient rule generally do depend on these variances. For rules 3 and 4, the losses are moderate when demand and inflation shocks are most important, but larger when the exchange rate shock is most important. That is, using r as the policy instrument

5. In the closed economy case, I continue to assume $\beta + \delta\theta = 1$. Therefore, since δ is zero, β is raised to one.

is most inefficient if there are large shocks to the r/e relation. In this case, r is an unreliable measure of the overall policy stance.

3.5 The Perils of Inflation Targeting

This section turns from instrument rules to target rules, specifically inflation targets. In the closed economy Svensson-Ball model, inflation targeting has good properties. In particular, the set of efficient Taylor rules is equivalent to the set of inflation target policies with different speeds of adjustment. In an open economy, however, inflation targeting can be dangerous.

3.5.1 Strict Inflation Targets

As in Ball (1997), strict inflation targeting is defined as the policy that minimizes the variance of inflation. When inflation deviates from its target, strict targeting eliminates the deviation as quickly as possible. I first evaluate this policy and then consider variations that allow slower adjustment.

Trivially, strict inflation targeting is an efficient policy: it minimizes the weighted sum of output and inflation variances when the output weight is zero. Strict targeting puts the economy at the northwest end of the variance frontier. In figure 3.1, the frontier is cut off when the output variance reaches 15; when the frontier is extended, the end is found at an output variance of 25.8 and inflation variance of 1.0. Choosing this point implies a huge sacrifice in output stability for a small gain in inflation stability. Moving down the frontier, the output variance could be reduced to 9.7 if the inflation variance were raised to 1.1, or to 4.1 if the inflation variance were raised to 1.6. Strict inflation targeting is highly suboptimal if policymakers put a nonnegligible weight on output.

The output variance of 25.8 compares to a variance of 8.3 under strict inflation targeting in the closed economy case. This difference arises from the different channels from policy to inflation. In a closed economy, the only channel is the one through output, which takes two periods (it takes a period for r to affect y and another period for y to affect π). With these lags, strict inflation targeting implies that policy sets expected inflation in two periods to zero. In an open economy, by contrast, policy can affect inflation in one period through the direct exchange rate channel. When policymakers minimize the variance of inflation, they set next period's expected inflation to zero:

$$E\pi_{+1} = 0. \tag{9}$$

Equation (9) implies large fluctuations in the exchange rate because next period's inflation can be controlled only by this variable. Intuitively, inflation in domestic goods prices cannot be influenced in one period, so large shifts in import prices are needed to move the average price level. (The appendix formalizes this interpretation.) The large shifts in exchange rates cause large output fluctuations through the IS curve.

This point can be illustrated with impulse response functions. Substituting equation (5) into equation (9) yields the instrument rule implied by strict inflation targeting:

$$(10) \qquad e = (\alpha/\gamma)y + (1/\gamma)(\pi + \gamma e_{-1}).$$

(Note this is a limiting case of eq. [7] in which the MCI equals the exchange rate.) Using equations (4), (5), and (10), I derive the dynamic effects of a unit shock to the Phillips curve. Figure 3.2 presents the results. Inflation returns to target after one period, but the shock triggers oscillations in the exchange rate and output. The oscillations arise because the exchange rate must be shifted each period to offset the inflationary or deflationary effects of previous shifts. These results contrast to strict inflation targeting in a closed economy, where an inflationary shock produces only a one-time output loss.[6]

3.5.2 The Case of New Zealand

These results appear to capture real-world experiences with inflation targeting, particularly New Zealand's pioneering policy in the early 1990s. During that period, observers criticized the Reserve Bank for moving the exchange rate too aggressively to control inflation. For example, Dickens (1996) argues that "whiplashing" of the exchange rate produced instability in output. He shows that aggregate inflation was steady because movements in import inflation offset movements in domestic goods inflation. These outcomes are similar to the effects of inflation targeting in my model.

Recently, the Reserve Bank has acknowledged problems with strict inflation targeting:

> If the focus of policy is limited to a fairly short horizon of around six to twelve months, the setting of the policy stance will tend to be dominated by the relatively rapid-acting direct effects of exchange rate and interest rate changes on inflation. In the early years of inflation targeting, this was, in fact, more or less the way in which policy was run. . . . Basing the stance of policy solely on its direct impact on inflation, however, is hazardous. . . . It is possible that in some situations actions aimed at maintaining price stability in the short term could prove destabilizing to activity and inflation in the medium term. (Reserve Bank of New Zealand 1996, 28–29)

The Reserve Bank's story is similar to mine: moving inflation to target quickly requires strong reliance on the direct exchange rate channel, which has adverse side effects on output.[7]

6. Black, Macklem, and Rose (1997) find that strict inflation targeting produces a large output variance in simulations of the Bank of Canada's model. Their interpretation of this result is similar to mine.

7. The Reserve Bank discusses direct inflation effects of interest rates as well as exchange rates because mortgage payments enter New Zealand's consumer price index.

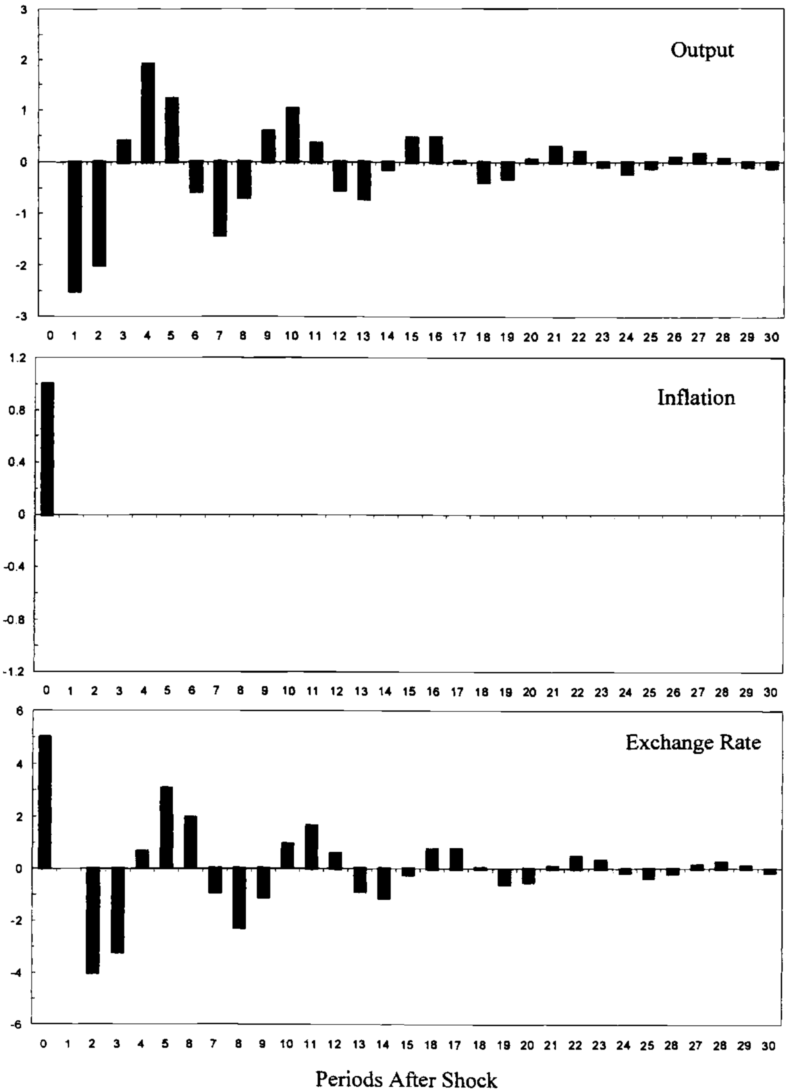

Fig. 3.2 Strict inflation targets: responses to an inflation shock

3.5.3 Gradual Adjustment?

The problems with strict inflation targeting have led observers to suggest a modification: policy should move inflation to its target more slowly. The Reserve Bank of New Zealand has accepted this idea; it reports that "in recent years the Bank's policy horizon has lengthened further into the future" and that this means it relies more heavily on the output channel to control inflation (1996, 29).

In the current model, however, it is not obvious what policy rule captures the goal of "lengthening the policy horizon." One natural idea (suggested by several readers) is to target inflation two periods ahead rather than one period:

$$E\pi_{+2} = 0. \tag{11}$$

This condition is the one implied by strict inflation targeting in the closed economy model. In that model, the condition does not produce oscillations in output.

In the current model, however, equation (11) does not determine a unique policy rule. Since policy can control inflation period by period, there are multiple paths to zero inflation in two periods. By the law of iterated expectations, $E\pi_{+1} = 0$ in all periods implies $E\pi_{+2} = 0$ in all periods. Thus a strict inflation target is one policy that satisfies equation (11). But there are other policies that return inflation to zero in two periods but not one period.[8]

The same point applies to various modifications of equation (11). For example, in the closed economy model, any efficient policy can be written as an inflation target with slow adjustment: $E\pi_{+2} = qE\pi_{+1}$, $0 \leq q \leq 1$. This condition is also consistent with multiple policies in the current model. There does not appear to be any simple restriction on inflation that implies a unique policy with desirable properties. Policymakers who wish to return inflation to target over the "medium term" need some additional criterion to define their rule.[9]

3.6 Long-Run Inflation Targets

This section presents the good news about inflation targets. The problems described in the previous section can be overcome by modifying the target variable. In light of earlier results, a natural modification is to target long-run inflation, π^*.

3.6.1 The Policies

Strict long-run inflation targeting is defined as the policy that minimizes the variance of $\pi^* = \pi + \gamma e_{-1}$. To see its implications, note that equation (2) can be rewritten as

$$\pi^* = \pi^*_{-1} + \alpha y_{-1} + \eta. \tag{12}$$

8. An example is a rule in which policy makes no contemporaneous response to shocks, but the exchange rate is adjusted after one period to return inflation to target in two periods.

9. Another possible rule is partial adjustment in one period: $E\pi_{+1} = q\pi$. This condition defines a unique policy, but the variance of output is large. The condition implies the same responses to demand and exchange rate shocks as does strict inflation targeting. These shocks have no contemporaneous effects on inflation, so policy must fully eliminate their effects in the next period, even for $q > 0$.

This equation is the same as a closed economy Phillips curve, except that π^* replaces π. The exchange rate is eliminated, so policy affects π^* only through the output channel. Thus policy affects π^* with a two-period lag, and strict targeting implies

$$E\pi^*_{+2} = 0. \tag{13}$$

In contrast to a two-period-ahead target for total inflation, equation (13) defines a unique policy.

There are two related motivations for targeting π^* rather than π. First, since π^* is not influenced by the exchange rate, policy uses only the output channel to control inflation. This avoids the exchange rate "whiplashing" discussed in the previous section. Second, as discussed in section 3.3, π^* gives the level of inflation with transitory exchange rate effects removed. Targeting π^* keeps underlying inflation on track.

In addition to strict π^* targeting, I consider gradual adjustment of π^*:

$$E\pi^*_{+2} = qE\pi^*_{+1}, \qquad 0 \leq q \leq 1. \tag{14}$$

This rule is similar to the gradual adjustment rule that is optimal in a closed economy. Policy adjusts $E\pi^*_{+2}$ part of the way to the target from $E\pi^*_{+1}$, which it takes as given. The motivation for adjusting slowly is to smooth the path of output.

In practice, countries with inflation targets do not formally adjust for exchange rates in the way suggested here. However, adjustments may occur implicitly. For example, a central bank economist once told me that inflation was below his country's target but that this was desirable because the currency was temporarily strong and policy needed to "leave room" for the effects of depreciation. Keeping inflation below its official target when the exchange rate is strong is similar to targeting π^*.

3.6.2 Results

To examine π^* targets formally, I substitute equations (12) and (1) into condition (14). This leads to the instrument rule implied by π^* targets:

$$w'r + (1 - w')e = a'y + b'\pi^*, \tag{15}$$

where

$$w' = \beta/(\beta + \delta), \quad a' = (1 - q + \lambda)/(\beta + \delta), \quad b' = (1 - q)/[\alpha(\beta + \delta)].$$

This equation includes the same variables as the optimal rule in section 3.3, but the coefficients are different. The MCI weights are given exactly by the relative sizes of β and δ; for base parameters, $w' = 0.75$. The coefficients on y and π^* depend on the adjustment speed q.

The appendix calculates the variances of output and inflation under π^* targeting. Figure 3.3 plots the results for q between zero and one. The case of

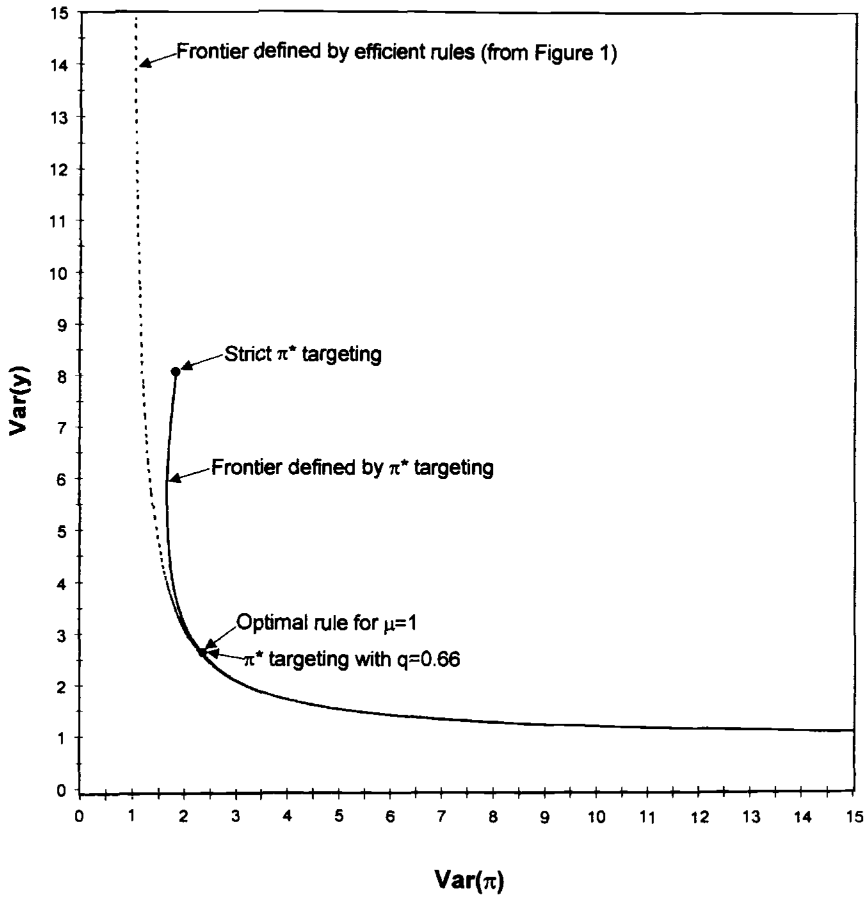

Fig. 3.3 π^* Targeting

strict π^* targeting corresponds to the northwest corner of the curve. For comparison, figure 3.3 also plots the set of efficient policies from figure 3.1.

The figure shows that targeting π^* produces more stable output than targeting π. This is true even for strict π^* targets, which produce an output variance of 8.3, compared to 25.8 for π targets. Figure 3.4 shows the dynamic effects of an inflation shock under π^* targets and confirms that this policy avoids oscillations in output. Strict π^* targeting is, however, moderately inefficient. There is an efficient instrument rule that produces an output variance of 8.3 and an inflation variance of 1.2. Strict π^* targets produce the same output variance with an inflation variance of 1.9.

As the parameter q is raised, so adjustment becomes slower, we move southeast on the frontier defined by π^* targeting. This frontier quickly moves close

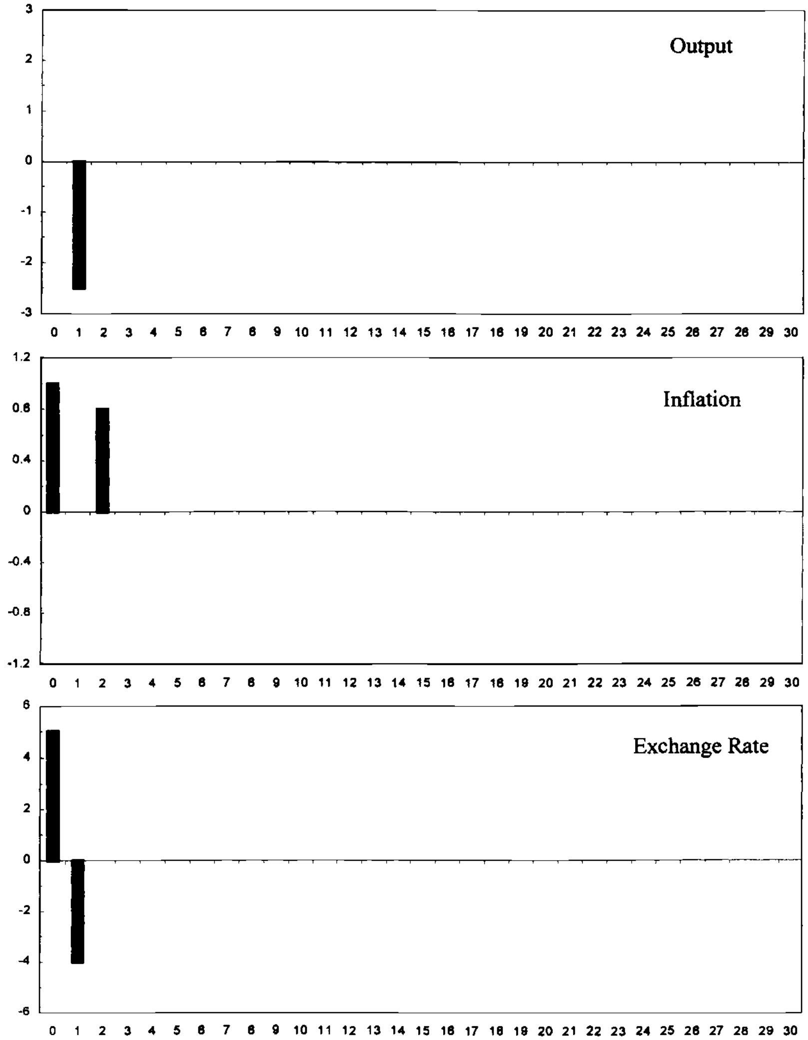

Fig. 3.4 Strict π^* targets: responses to an inflation shock

to the efficient frontier. Thus, as long as policymakers put a nonnegligible weight on output variance, there is a version of π^* targeting that closely approximates the optimal policy. For example, for equal weights on inflation and output variances, the optimal policy has an MCI weight w of 0.70 and output and π^* coefficients of 1.35 and 1.06. For a π^* target with q = 0.66, the corresponding numbers are 0.75, 1.43, and 1.08. The variances of output and inflation are 2.50 and 2.44 under the optimal policy and 2.48 and 2.48 under π^* targeting.

3.7 Conclusion

In a closed economy, inflation targeting and Taylor rules perform well in stabilizing both output and inflation. In an open economy, however, these policies perform poorly unless they are modified. Specifically, if policymakers minimize a weighted sum of output and inflation variances, their policy instrument should be an MCI based on both the interest rate and the exchange rate. The weight on the exchange rate is equal to or slightly greater than this variable's relative effect on spending. As a target variable, policymakers with this paper's objective function should choose "long-run inflation"—an inflation variable purged of the transitory effects of exchange rate fluctuations. This variable should also replace inflation on the right-hand side of the instrument rule.

Several countries currently use an MCI as their policy instrument. In addition, some appear to have moved informally toward targeting long-run inflation, for example by keeping inflation below target when a depreciation is expected. A possible strengthening of this policy is to make long-run inflation the formal target variable. In practice, this could be done by adding an adjustment to calculations of "underlying" inflation: the effects of the exchange rate could be removed along with other transitory influences on inflation. At least one private firm in New Zealand already produces an underlying inflation series along these lines (Dickens 1996).

Appendix

Domestic Goods and Imports

Here I derive the Phillips curve, equation (2), from assumptions about inflation in the prices of domestic goods and imports. Domestic goods inflation is given by

$$\pi^{d} = \pi_{-1} + \alpha' y_{-1} + \eta'. \tag{A1}$$

This equation is similar to a closed economy Phillips curve: π^{d} is determined by lagged inflation and lagged output.

To determine import price inflation, I assume that foreign firms desire constant real prices in their home currencies. This implies that their desired real prices in local currency are $-e$. However, they adjust their prices to changes in e with a one-period lag. Like domestic firms, they also adjust their prices based on lagged inflation. Thus import inflation is

(A2) $$\pi^{m} = \pi_{-1} - (e_{-1} - e_{-2}).$$

Finally, aggregate inflation is the average of (A1) and (A2) weighted by the shares of imports and domestic goods in the price index. If the import share is γ, this yields equation (2) with $\alpha = (1 - \gamma)\alpha'$ and $\eta = (1 - \gamma)\eta'$.

The Variances of Output and Inflation

Here I describe the computation of the variances of output and inflation under alternative policies. Consider first the rule given by equations (6) and (7). Substituting equations (4) and (5) into equation (6) yields an expression for the exchange rate in terms of lagged e, π, and y

(A3) $$e = (\lambda z + \alpha n)y_{-1} + n\pi_{-1} - (\beta/\theta + \delta)ze_{-1} + \gamma ne_{-2} + z\varepsilon + n\eta + \beta m\nu + \beta z\nu_{-1}, \qquad z = \lambda m + \alpha n.$$

This equation and equations (4) and (5) define a vector process for e, π, and y:

(A4) $$X = \Phi_1 X_{-1} + \Phi_2 X_{-2} + E,$$

where $X = [y\ \pi\ e]'$.

The elements of E depend on the current and once-lagged values of white noise shocks. Thus E follows a vector MA(1) process with parameters determined by the underlying parameters of the model. X follows an ARMA(2, 1) process. For given parameter values and given values of the constants m and n, one can numerically derive the variance of X using standard formulas (see Hendry 1995, sec. 11.3). To determine the set of efficient policies, I search over m and n to find combinations that minimize a weighted sum of the output and inflation variances.

To determine the variances of output and inflation under a π^* target, note that equation (15) is equivalent to equation (7) with m set to $\theta/(\theta\delta + \beta)$ and n set to $\theta(1 - q)/[\alpha(\theta\delta + \beta)]$. For a given q, the variances of output and inflation under equation (15) are given by the variances for the equivalent version of equation (7).

References

Ball, Laurence. 1997. Efficient rules for monetary policy. NBER Working Paper no. 5952. Cambridge, Mass.: National Bureau of Economic Research, March.

Black, Richard, Tiff Macklem, and David Rose. 1997. On policy rules for price stability. In *Price stability, inflation targets and monetary policy*, ed. Bank of Canada. Ottawa: Bank of Canada.

Dickens, Rodney. 1996. Proposed modifications to monetary policy. Wellington: Ord Minnett New Zealand Research, May.

Duguay, Pierre. 1994. Empirical evidence on the strength of the monetary transmission mechanism in Canada: An aggregate approach. *Journal of Monetary Economics* 33 (February): 39–61.
Freedman, Charles. 1994. The use of indicators and of the monetary conditions index in Canada. In *Frameworks for monetary stability,* ed. Tomas J. T. Balino and Carlo Cottarelli. Washington, D.C.: International Monetary Fund.
Gerlach, Stefan, and Franks Smets. 1996. MCIs and monetary policy in small open economies under floating exchange rates. Basel: Bank for International Settlements, November.
Gruen, David, and Geoffrey Shuetrim. 1994. Internationalisation and the macroeconomy. In *International integration of the Australian economy,* ed. Philip Lowe and Jacqueline Dwyer. Sydney: Reserve Bank of Australia.
Hendry, David F. 1995. *Dynamic econometrics.* Oxford: Oxford University Press.
Lafleche, Therese. 1996. The impact of exchange rate movements on consumer prices. *Bank of Canada Review,* winter 1996–97, 21–32.
Longworth, David J., and Stephen S. Poloz. 1986. A comparison of alternative monetary policy regimes in a small dynamic open-economy simulation model. Bank of Canada Technical Report no. 42. Ottawa: Bank of Canada, April.
Reserve Bank of New Zealand. 1996. *Briefing on the Reserve Bank of New Zealand.* Wellington: Reserve Bank of New Zealand, October.
Svensson, Lars E. O. 1997. Inflation forecast targeting: Implementing and monitoring inflation targets. *European Economic Review* 41:1111–46.
———. Forthcoming. Open-economy inflation targeting. *Journal of International Economics.*
Taylor, John B. 1994. The inflation/output variability trade-off revisited. In *Goals, guidelines, and constraints facing monetary policymakers,* ed. Jeffrey C. Fuhrer. Boston: Federal Reserve Bank of Boston.

Comment Thomas J. Sargent

Summary

I offer two related, apparently but not necessarily contradictory, criticisms of Laurence Ball's model. First, the model incorporates too little rational expectations; second, it assumes too much rational expectations. The private sector uses too little rational expectations (at least compared to the government); and the government uses too much (in view of how little rational expectations Ball has built into the private sector's model—and in view of Ball's empirical approach to his model as an approximation). Let me explain from the beginning.

Ball ascribes a return function to the government (a weighted sum of unconditional variances of inflation and output) then solves a "modern" single-agent control problem to deduce an optimal rule for monetary policy. He models the economy as being like *nature,* a time-invariant system of stochastic difference

Thomas J. Sargent is senior fellow at the Hoover Institution, the Donald Lucas Professor of Economics at Stanford University, and a research associate of the National Bureau of Economic Research.

equations, driven by serially uncorrelated shocks, presenting state and control vectors to the policymaker.[1] The policymaker views the model economy as "known" and "fixed" (with respect to the policymaker's choices). Though Ball's model is estimable using modern methods (the "duals" of the control theoretic methods he uses), Ball abstains from estimation and instead resorts to "calibration." In calibrating, Ball implicitly treats his model somehow as an approximation, not to be taken literally empirically. This empirical approach is compatible with Ball's motivation of his work as a sensitivity exercise designed to explore how augmenting the government's model to incorporate aspects of a foreign sector will affect the character of optimal rules that he had deduced earlier from a related closed economy model (Ball 1997). Ball does a good job of explaining the alterations he has made to the baseline closed economy model, of interpreting how they make the optimal rule change, and of informally explaining how the optimal rule corresponds to practices observed in some small open economies. The optimal rule contains the "advice" Ball coaxes from his formal model.[2]

Ball also evaluates six arbitrary rules. Within the confines of this single paper, comparing the operation of these rules is of less interest than just looking at the optimal rule. These comparisons acquire interest as ingredients of a robustness calculation only when they are put together with the workings of these six rules within the *different* models represented in the other papers in the conference.

To supplement Ball's acknowledgment in the text that the government's model is "pre–Lucas critique" (although cleverly constructed to embody one version of the natural rate hypothesis), I recommend Robert King and Mark Watson's (1994) "revisionist history" of the Phillips curve. King and Watson use theory and empirical evidence to document how the type of "unit root" specification of the Phillips curve used by Ball reflects the serial correlation structure of inflation detected only over the last 25 years or so, which emerged hand in hand with our current fiat monetary policy regime. Our earlier monetary regime was associated with much lower serial correlation of inflation, and a different apparent intertemporal Phillips curve.[3] I recommend Ball's (1995) *Journal of Monetary Economics* paper (where the government sets a repeated economy strategy and the private economy makes predictions that reflect the government's strategy) as an example of an analysis with optimizing behavior on all sides, a setup not subject to my first criticism (though possibly to my second one).

1. Ball thus adheres to Einstein's dictum that while nature might be complicated, it is not cruel. Contrast this modern view with the malevolent view of nature embraced by the "robust" monetary authority to be described shortly.

2. He computes an optimal rule for each value of a relative weight parameterizing the government's objective function.

3. Robert Hall makes related and more extensive observations in his comment on chap. 9 of this volume.

Ball can argue somewhat convincingly that despite imperfect "microfoundations," and despite vulnerability to the Lucas critique, a model of his type can somehow be a good approximation to a "truer" (bigger? general equilibrium?) model. Accepting this defense against the charge of too little rational expectations in the model strengthens my second criticism and prompts retreat from the rational expectations that Ball ascribes to the government in formulating its control problem. If the model is an approximation, the optimal policy rule ought to incorporate this attitude. This attitude recommends "postmodern" or "robust" rather than "modern" control theory. A quest for robustness overshadows this conference, so I want to explore how a preference for robustness would affect Ball's analysis.

Modern and Robust Government Problems

Ball's control problem is a member of the following class. Where x_t and u_t are the state and control vectors and z_t is a vector representing the decision maker's rewards, the government's model is

(1a) $$x_{t+1} = Ax_t + Bu_t + Cw_{t+1},$$

(1b) $$u_t = -Fx_t,$$

(1c) $$z_t = Hx_t,$$

where $\{w_{t+1}\}$ is a martingale difference sequence with unity contemporaneous covariance matrix, adapted to its own history.[4] The government chooses a feedback rule of the form (1b) to achieve an objective that Ball formulates as the unconditional variance of z_t. For the purpose of formalizing the control problem, it is useful to express the model in the form

(2) $$z_t = G(L)w_t \equiv H[I - (A - BF)L]^{-1}Cw_t,$$

Equation (2) shows how the white noise shocks w_t affect the variables the government cares about. For Ball, the problem is to maximize minus the unconditional variance of z, by choice of F. This can be expressed in the frequency domain as choosing an F to maximize

4. For Ball's model, we can take

$$x_t = \begin{bmatrix} y_t & \pi_t & e_{t-1} & v_t \end{bmatrix}', \quad u_t = e_t, \quad A = \begin{bmatrix} \lambda & 0 & 0 & \beta/\theta \\ \alpha & 1 & \gamma & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix},$$

$$B' = [-(\beta/\theta + \delta) \quad -\gamma \quad 1 \quad 0], \quad C = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & \sqrt{2} \end{bmatrix}.$$

(3) $$H^2 = -Ez_t z_t' = -\int_\Gamma [G(\zeta)G(\zeta)']\, d\zeta,$$

where ζ is a complex variable, Γ is the unit circle, and the prime denotes matrix transposition and complex conjugation.

This problem is "modern" in the sense that the government is assumed to know enough about the model to have rational expectations. It treats the model as known and true. This means that the government knows: (1) $G(L)$ and how it depends on F and (2) that the shock process w_t is serially uncorrelated.

Robust control theory would back away from the rational expectations assumption and endow the government with the view that while "good," its model is not "true," only an approximation. Suspicion that the model is an approximation imparts a preference for "robustness," that is, acceptable performance of a rule across a range of models "in the vicinity" of the government's—"in the vicinity" because the model is viewed as a useful approximation. To formalize this approach, robust decision theory treats the model not as true but as a base or reference model around which the government suspects approximation errors.[5] The key is to describe the form that the approximation errors can take, and how "big" they can be. Approximation errors have fruitfully been formulated as showing up as a shock process $\{w_t\}$ that is arbitrarily serially correlated, not serially uncorrelated as specified. If the shocks have spectral density $S_w(\zeta)$ rather than I, then the variance of z_t is not given by the second equality in equation (3) but by

(4) $$\tilde{H}_2 = -\int_\Gamma G(\zeta)S_w(\zeta)G(\zeta)'\, d\zeta.$$

The robust decision maker seeks a rule that delivers an "acceptable" $\tilde{H}_2$ over a domain of somewhat arbitrary $S_w(\zeta)$. To construct the domain of approximation errors, some sort of "constant variance" restraint $\int_\Gamma S_w(\zeta)\, d\zeta = I$ is imposed on the potentially perturbed error processes.

To design a robust rule, the decision maker performs a worst case analysis. He solves a problem of the form

(5) $$\max_F \min_{S_w(\zeta)} \tilde{H}_2 \quad \text{such that} \quad \int_\Gamma S_w(\zeta)\, d\zeta = I.$$

The F that solves this max-min problem works better than the modern rule across much of the domain of unknown error serial correlation specifications. The associated minimizing $S_w(\zeta)$ is the "most pessimistic" view of the shock serial correlation process, an artifact of computing a rule designed to work well over a variety of $S_w(\zeta)$.[6]

5. For general background on robust control, see Zhou (1996), Başar and Bernhard (1995), and Mustafa and Glover (1990). See Whittle (1990) for an account of risk-sensitive control, and Hansen and Sargent (1995) for a recursive reformulation of risk sensitivity designed to accommodate discounting.

6. Max-min problems of this type occur widely in analyses of robustness. For a related problem in quite a different context, see Fudenberg and Levine (1995).

Notice how this formulation makes the specification error show up as a perturbed shock process that nevertheless feeds through the system in the way the model specifies. This is a restrictive way of modeling "misspecification," though as Hatanaka's (1975) work emphasized, arbitrarily serially correlated errors provide specifications so flexible that they often undermine econometric identifiability. A restrictive aspect of the present way of modeling specification errors is that the misspecified errors are supposed to feed through the system just as those in the reference model.

Remarkably, but maybe not surprisingly, because this *is* a form of Knightian uncertainty, there is a sense in which the robust decision maker turns back from the Lucas critique. Thus, for a *given* reference model, variations in the shock serial correlation process within the domain satisfying the constraint in problem (5) are presumed to leave the robust decision rule intact. Here I am interpreting the Lucas critique as but an application of the principle, already well reflected in "Keynesian" contributions like Kareken, Muench, and Wallace (1973),[7] stating that optimal decision rules depend on laws of motion or transition laws expressing the "constraints," which in the modern setting always include descriptions of the serial correlation patterns of the shocks. The robust controller's behavior partly belies that principle.[8]

Problem (5) embodies a game, not real, but a mental one, played in the mind of a monetary authority, who as an instrument for attaining a robust rule contemplates the reactions of a diabolical nature that is not "fixed" and that in response to *its* choice of *F*, responds by "choosing" a shock process serial correlation structure to harm the monetary authority. The point is not that the policymaker expects the worst but that by planning against it he assures acceptable performance under a range of specification errors.

Problem (5) leads to what is called the H^∞ formulation of robust decision theory.[9] A less extreme and more convenient formulation is the "minimum entropy" criterion, according to which the government just maximizes

$$H_{\text{entropy}} = -\log \det \int_\Gamma [\kappa I - G(\zeta)G(\zeta)'] \, d\zeta. \tag{6}$$

This criterion is defined only for $\kappa \geq \kappa_b$, where κ_b is the minimum positive scalar for which the integrand is positive semidefinite.[10] The decision rule F that maximizes criterion (6) approximates the decision rule F that solves prob-

7. Their monetary authority has rational expectations and recommends "look at everything" rules, just like Ball's.

8. The robust decision maker's behavior *does* vary systematically with respect to variations in his reference model, the version of the Lucas critique that survives under robust decision theory.

9. See Mustafa and Glover (1990); see Hansen and Sargent (1998) for the formulations of H^∞ control modified to incorporate discounting. Hansen and Sargent also describe the "minimum entropy" formulation below and how it relates to the H^∞ criterion. In an interesting example, Kasa (1998) studies how a robust controller would behave in ways that an outsider might interpret as responding to adjustment costs.

10. Hansen and Sargent (1998) derive this criterion for discounted problems relate it to the H^∞ and risk-sensitivity formulations.

Table 3C.1 **Robust Policy Rules**

σ	$\bar{a}$	$\bar{b}$
0	1.04	.82
−.1	1.17	1.13
−.2	1.55	2.08

lem (5) as $\kappa \searrow \kappa_b$. Further, the parameter $\kappa > 0$ is interpretable in terms of the risk-sensitive formulation of Whittle (1990) and Hansen and Sargent (1995). In particular, κ equals minus the reciprocal of the risk-sensitivity parameter, σ.[11] Formulation (6) provides a "smooth" one-parameter family of approximators of (5). For $\sigma = 0$, we have the modern case, while for σ approaching $-\kappa_b^{-1}$, we attain the case envisaged in (5).

Examples

Using criterion (6), I have computed decision rules for three values of $-\kappa^{-1} \equiv \sigma$, namely, $\sigma = 0, -.1$, and $-.2$. I prefer to parameterize in terms of σ because of its link to risk sensitivity, but this is just a matter of taste. The value $\sigma = 0$ corresponds to Ball's calculations and checks my calculations against his.[12] For Ball's parameter values and the case of equal weights on inflation and output variances in the government's objective, I calculated the responses of the interest rate, r to output y and the monetary conditions index π^* given in table 3C.1. Here $r = \bar{a}y + \bar{b}\pi^*$, where $\bar{a}$ and $\bar{b}$ are calculated just as in the $\sigma = 0$ case as explained by Ball. These numbers agree with Ball's in the $\sigma = 0$ case. Note how the interest rate becomes *more* sensitive to both y and π^* as the absolute value of σ rises (i.e., as the preference for robustness increases). This aspect of the rules illustrates that the "caution" induced by activating the preference for robustness does not translate into "doing less." Caution is relative to a worst case pattern of the shock process, which depends on the specification of the reference model. In particular, the specification of the model affects the magnitude and the serial correlation properties of the worst case shocks, and thus influences how the policymaker responds to actual shocks.[13]

To shed light on how a preference for robustness surfaces, it is useful to display the optimized $G(\zeta)$ functions in the frequency domain. Figure 3C.1 displays the maximum singular value of $G(\zeta)G(\zeta)'$ by frequency for $\sigma = 0$, $-.1$, and $-.2$. As σ rises in absolute value, the maximum singular value across

11. See Anderson, Hansen, and Sargent (1997) and Hansen and Sargent (1998). These papers also relate the risk-sensitive formulation to measures of model misspecification like those used by White (1982) and Sims (1972).

12. They agree where they should.

13. That robust control can be "more active" comes through in results reported by Zhou (1996, 438–39).

While the optimal rules remain linear, their dependence on innovation variances reflects a breakdown of "certainty equivalence."

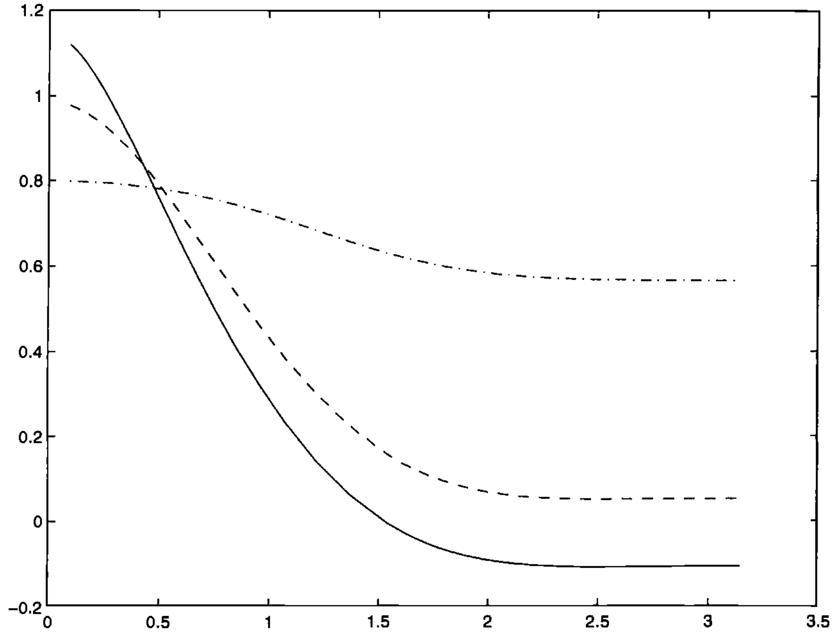

Fig. 3C.1 Logarithm of maximum singular values of $G(\zeta)$ for three values of σ

Note: $\sigma = 0, -.1$, and $-.2$ are indicated by solid, dashed, and dash-dotted lines, respectively. Angular frequency (from 0 to π) is the ordinate, the logarithm of the maximum singular value the coordinate.

frequencies falls, but the area under the curve rises.[14] This pattern captures how worst case performance improves as σ rises in absolute value, at the cost of worse performance at the reference model. In particular, think about how the "great deceiver" feared in problem (5) would respond to a $G(\zeta)$ having the pattern of maximum singular values associated with the $\sigma = 0$ case in figure 3C.1. The deceiver could lower the utility of the monetary authority (raise the weighted average of variances) by concentrating the spectral power of the shock process near the frequencies at which the maximum singular value is highest. To protect itself against such a malevolent nature, the monetary authority can design a rule to lower the maximum singular value, although this causes the average level of the maximum singular values across frequencies to drift upward (which lowers the value of the monetary authority's objective under the assumption that the reference model is *true*). In Ball's model, flattening the maximum singular values across frequencies can be achieved by making the interest rate respond more to both y and π^*.

14. The L^∞ norms for $\sigma = 0, -.1$, and $-.2$ are 3.13, 2.68, and 2.22; the L^2 norms are 1.57, 1.61, and 1.97.

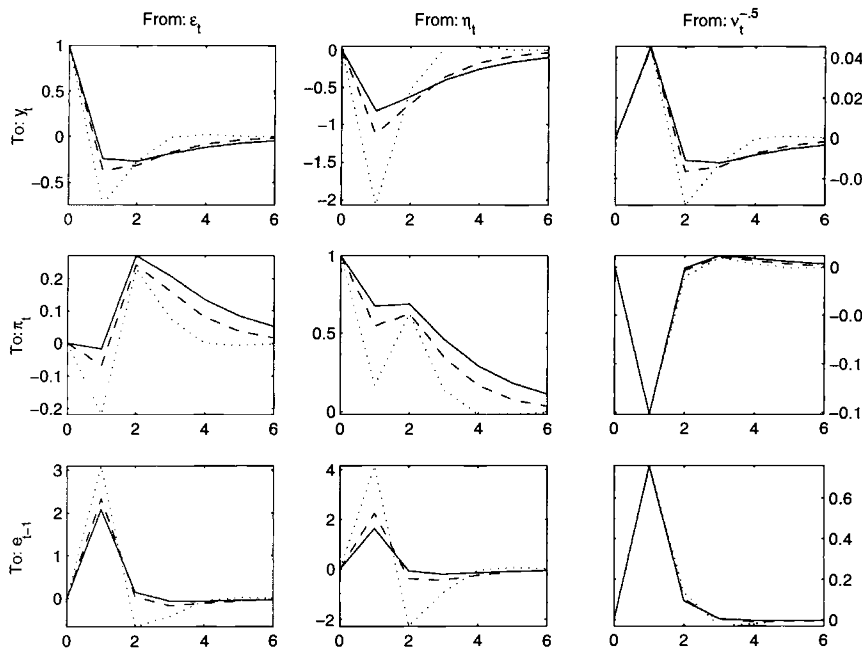

Fig. 3C.2 Impulse response of $[y\ \pi\ e_{-t}]'$ to $[\varepsilon_t\ \eta_t\ (1/\sqrt{2})v_t]'$ for Ball's model for three values of σ

Note: $\sigma = 0, -.1$, and $-.2$ are indicated by solid, dashed, and dotted lines. The shocks correspond to columns, the variables to rows.

Figure 3C.2 depicts the impulse response functions with respect to the w_ts (i.e., the $G(L)$s) for the three values of $\sigma = 0, -.1$, and $-.2$. Close inspection of these figures confirms how making σ more negative leads to "whiter" impulse responses. The "less colorful" impulse responses reflect the "flattening" across frequencies of the maximum singular values as σ (which is nonpositive) is raised in absolute value. Thus the "more active" responses in the rules for more negative σs, depicted in table 3C.1, deliver whiter spectral outcomes (assuming the reference model to be true).

Thus this particular way of framing the class of misspecifications to be protected against causes "more caution" to translate into "more activity." How can this happen? One way to approach this question is to study how the "worst case" conditional means of the shocks, call them $\hat{w}_{t+j}$, respond to the true shocks w_t. Hansen, Sargent, and Tallarini (forthcoming) report a formula for such worse case conditional means. The formula shows how the worst case conditional means are linear functions of the state of the system x_t. This makes it possible to compute impulse responses of the worst case shocks with respect to the w_ts. Figure 3C.3 depicts the impulse responses of these worst case conditional means in the shocks with respect to the w_ts for $\sigma = -.1$ and $-.2$; for

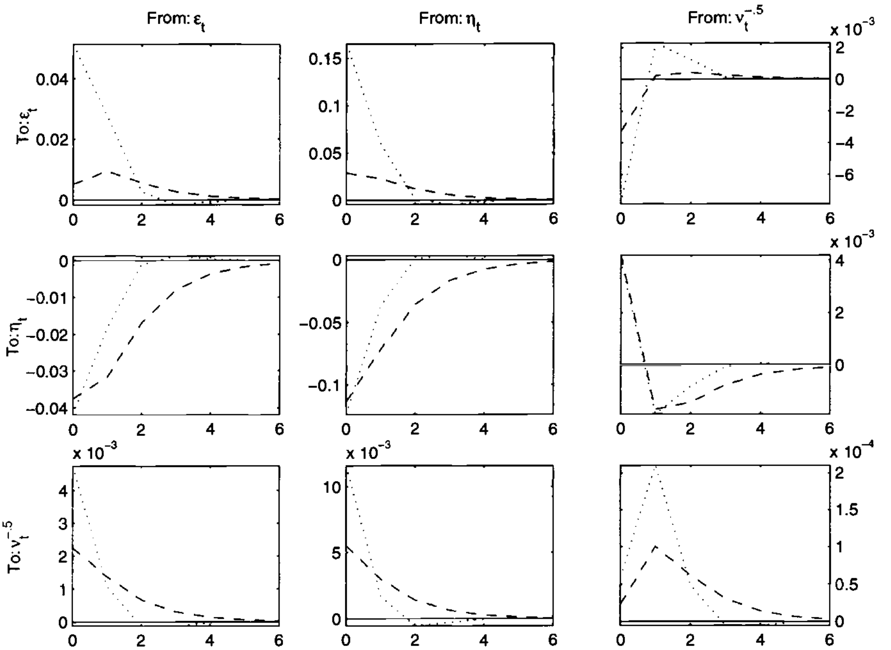

Fig. 3C.3 Impulse response of the worst case $[\varepsilon_t\ \eta_t\ (1/\sqrt{2})\nu_t]'$ with respect to $[\varepsilon_t\ \eta_t\ (1/\sqrt{2})\nu_t]'$ for Ball's model for two values of σ

Note: $\sigma = -.1$ and $-.2$ are indicated by dashed and dotted lines. The shocks correspond to columns, the worst case shocks to rows.

$\sigma = 0$, these worst case means are identically zero. The impulse responses in figure 3C.3 show how for Ball's model, the worst case conditional means in the shocks change both in size and in their serial correlation properties as σ becomes more negative. Figure 3C.3 shows how the worst case conditional means of shocks become larger as σ grows in absolute value, and that they respond as though the worst case shocks are serially correlated.[15] This makes the robust decision maker respond to what the reference model asserts are white noise shocks as though they were positively serially correlated. To guard against the worst case, the robust decision maker exercises caution by "pressing the brakes" or "pushing on the accelerator" more than would the modern decision maker. This is reasonable because a shock is interpreted as being larger and more persistent than it is supposed to be by the reference model. This helps to explain the pattern of the table 3C.1 responses.

15. The figures also show that the worst case shocks are less serially correlated for the σ that is larger in absolute value. In interpreting this diminished serial correlation, it has to be remembered that these objects are the outcome of the "game" associated with a minimax problem, and that the great deceiver picking these worst case errors faces a different feedback rule F for each σ.

Thus, in the context of Ball's model, the caution exercised by the robust decision maker does not manifest itself in the form of moving less in response to shocks. Robust decision theory teaches that what caution means depends on the model with respect to which caution is being exercised.

Unanswered Questions

The preceding response to my two criticisms of Ball's general approach opens more questions than it answers. If we imagine the monetary authority to be using robust decision theory as I have outlined, what should be done to close the model, in the spirit of seeking some counterpart to a rational expectations equilibrium, only now where both the government and the market are robust decision makers? More generally, if policymakers are entertaining the same doubts across models as we are at this conference, how are we to formulate the forecasting problems facing private agents? Along with the papers presented in this conference, my "robust decision" analysis of Ball's model proceeds without any sort of "robust equilibrium" concept.[16] Maybe such an equilibrium concept could help us.[17]

References

Anderson, Evan. 1998. Uncertainty and the dynamics of Pareto optimal allocations. Chicago: University of Chicago. Manuscript.

Anderson, Evan, L. P. Hansen, and T. J. Sargent. 1997. Robustness and risk-sensitivity in general equilibrium. Chicago and Palo Alto: University of Chicago and Hoover Institution. Manuscript.

Ball, Laurence. 1995. Time-consistent policy and persistent changes in inflation. *Journal of Monetary Economics* 36 (2): 329–50.

———. 1997. Efficient rules for monetary policy. NBER Working Paper no. 5952. Cambridge, Mass.: National Bureau of Economic Research.

Başar, Tamar, and Pierre Bernhard. 1995. *H^∞-optimal control and related minimax design problems: A dynamic game approach.* Boston: Birkhäuser.

Fudenberg, Drew, and David Levine. 1995. Universal consistency and cautious fictious play. *Journal of Economic Dynamics and Control* 19 (July): 1065–89.

Hansen, L. P., and T. J. Sargent. 1995. Discounted linear exponential quadratic Gaussian control. *IEEE Transactions on Automatic Control* 40 (May): 968–71.

———. 1998. Alternative representations of discounted robust linear quadratic control. Chicago and Palo Alto: University of Chicago and Hoover Institution. Manuscript.

16. E.g., the agents inside the four models compared in chap. 6 of this volume, by Levin, Wieland, and Williams, all live in secure rational expectations environments and have no need to trouble themselves with the doubts over models that preoccupy Levin, Wieland, and Williams. I believe that the three authors' comforting finding of robustness across models might mainly reflect the proximity of the four models. Note that the four models compared are all rational expectations models with "forward-looking" private sectors, and all have been thoughtfully designed to fit recent U.S. data and to incorporate similar monetary control channels. It would be more challenging for rules also to be robust against the non-forward-looking models illustrated by those of Ball and, in chap. 5, of Rudebusch and Svensson.

17. Evan Anderson (1998) studies a multiple-agent environment with robust decision makers whose preferences for robustness diverge.

Hansen, L. P., T. J. Sargent, and T. D. Tallarini Jr. Forthcoming. Robust permanent income and pricing. *Review of Economic Studies.*
Hatanaka, Michio. 1975. On the global identification of the dynamic simultaneous equations model with stationary disturbances. *International Economic Review* 16 (3): 545–54.
Kareken, J. A., T. Muench, and N. Wallace. 1973. Optimal open market strategy: The use of information variables. *American Economic Review* 68 (1): 156–72.
Kasa, Kenneth. 1998. Observational equivalence of H^{∞} control problems. San Francisco: Federal Reserve Bank of San Francisco. Manuscript.
King, Robert G., and Mark W. Watson. 1994. The post-war U.S. Phillips curve: A revisionist econometric history. *Carnegie-Rochester Conference Series on Public Policy* 41:157–219.
Mustafa, Denis, and Keith Glover. 1990. *Minimum entropy H^{∞} control.* Berlin: Springer.
Sims, C. A. 1972. Approximate prior restrictions on distributed lag estimation. *Journal of the American Statistical Association* 67 (1): 169–75.
White, Halbert. 1982. Maximum likelihood estimation of misspecified models. *Econometrica* 50 (1): 1–25.
Whittle, Peter. 1990. *Risk-sensitive optimal control.* New York: Wiley.
Zhou, Kemin, with John C. Doyle and Keith Glover. 1996. *Robust and optimal control.* London: Prentice-Hall.

Discussion Summary

Regarding Sargent's discussion of the paper, *Bob Hall* conjectured that since a Bayesian does not respect the Knightian distinction between risk and uncertainty, all knowledge could be properly summarized in a Bayesian framework. *Sargent* replied that with an infinite-dimensional parameter space as given here, the Bayes consistency theorem breaks down. As an early reference on this issue, Sargent cited Christopher Sims's thesis.

David Longworth noted that the model's policy recommendation was very similar to what the Bank of Canada has been doing in the past, for example, deriving the MCI weights from the IS curve and carefully considering the exchange rate effect on inflation. A simple extension of the model would be to allow for cross-correlation between the error in the IS curve and the error in the exchange rate equation. The reason this is interesting is that this correlation, that is, the degree to which exchange rates are driven by nonfinancial fundamentals, determines the desirability of the MCI as a target variable in the short run. For example, the Reserve Bank of Australia and the Bank of England both tend to think that there is a high correlation between these error terms for their respective home countries. Under those circumstances, it is not desirable to leave the MCI unchanged over one period. On the contrary, the Bank of Canada believes this correlation to be very small in Canada (especially when one abstracts from shocks to commodity prices and from readily observable shocks to the IS curve).

Laurence Meyer remarked that this is a timely paper for two reasons. First, exchange rates have been very important through import prices in the recent U.S. inflation experience. Second, globalization leads to increased openness that affects the way monetary policy should respond. Meyer then expressed reluctance to use the MCI as a policy instrument. He argued that the optimal response of interest rates to changes in the exchange rate depends on the source and the persistence of the change. Also, the MCI is an incomplete measure of overall financial conditions. It leaves out such variables as long interest rates and equity prices, which might be equally or more important. Meyer proposed interpreting the paper as having an interest rate instrument that reacts to the exchange rate.

Frederic Mishkin mentioned the New Zealand experience, where exchange rate effects are much faster than interest rate effects. That creates an instrument instability and controllability problem as described in the paper, in particular with rigid bands and a short horizon.

Andrew Levin noted that the appendix of his paper with Wieland and Williams presented at the conference summarized some analysis of the importance of openness for U.S. policy rules. Stochastic simulations of the Taylor multicountry model and deterministic simulations of the Federal Reserve Board global model show that the optimal policy rules are essentially identical for different assumptions on foreign policy. This can be interpreted as evidence for the size of the external sector in the U.S. economy being small and the passthrough of exchange rates into prices being slow. In the context of Ball's model, this means that the δ and γ coefficients seem to be very small for the United States. *Ball* responded that the applicability of the model to the case of the United States is still an open question. He also expressed surprise over the ongoing discussion about the importance of recent events in Asia for the United States.

Lars Svensson noted that the Bank of Sweden is not using the MCI but a repurchase rate as a policy instrument. The MCI is used as a measure of the impact of monetary policy on aggregate demand.

Svensson also remarked that the exchange rate equation, equation (3) in the paper, is not consistent with the exchange rate being an asset price, that is, with interest rate parity. Furthermore, Svensson had recently addressed this problem in a forward-looking framework where he obtained an MCI that depended on the real exchange rate and a long real interest rate, as opposed to the short rate. Thus the resulting MCI in such a framework is very different from the measure being used by actual central banks. *Richard Clarida* suggested interpreting equation (3) as a link between the log real exchange rate and an interest differential. Such an equation can be derived from underlying principles as in Campbell and Clarida (1987), but a necessary condition for that proportionality to hold is that the real interest rate follow an AR(1) process. Without this assumption, the entire term structure of interest rates enters the equation. Moreover, when the policymaker chooses the exchange rate e, he only chooses the current interest rate r in the paper. In a more general model, this amounts to choosing the entire time path of r, so that the resulting MCI would not have as nice a representation in terms of current variables as in the paper.

Ben McCallum wondered about the paper's use of real variables to measure the MCI. First, the MCI used by actual central banks like New Zealand and Canada is measured in terms of nominal variables. How should such an MCI measure be deflated? This is not clear since subtracting the expected inflation rate would be appropriate to calculate real interest rates and dividing by the price level would be appropriate to obtain the real exchange rate. Second, policy indicators should ideally be nominal variables instead of real ones.

John Taylor noted that small economies have a desire to prevent exchange rate fluctuations and many papers at the conference discuss interest rate smoothing. What about exchange rate smoothing? *Lars Svensson* answered that with a short horizon for inflation targeting, there is a tendency to move the exchange rate too much. With a longer horizon, the exchange rate is smoothed to stabilize the consumer price index. *David Longworth* added that it would be important to include tables in the paper showing the variance of interest and exchange rates across different policy rules. This is particularly important for the comparison of rules targeting π (actual inflation) and π^* (long-run inflation).

Nicoletta Batini remarked that it would be interesting to see how rules with different target specifications, for example domestic inflation, compare to the "long-run" inflation target in the paper. Use of targets that—in an economy that imports only final goods—are not influenced by the exchange-rate-to-import-price channel may remove the need for an ad hoc target à la Ball. Also, the use of an MCI as an instrument is not desirable because it requires measuring the equilibrium exchange rate in addition to the equilibrium interest rate and the output gap as in more conventional, Taylor-type simple rules.

Reference

Campbell, John, and Richard Clarida. 1987. The dollar and real interest rates: An empirical investigation. *Carnegie-Rochester Conference Series on Public Policy* 27: 103–39.