International Finance 2:1, 1999: pp. 63–83

Efficient Rules for Monetary Policy*

Laurence Ball
Johns Hopkins University, Baltimore.

Abstract

This paper defines an efficient rule for monetary policy as one that minimizes a weighted sum of output variance and inflation variance. It derives several results about the efficiency of alternative rules in a simple macroeconomic model. First, efficient rules can be expressed as 'Taylor rules' in which interest rates respond to output and inflation. But the coefficients in efficient Taylor rules differ from the coefficients that fit actual policy in the United States. Second, inflation targets are efficient. Indeed, the set of efficient rules is equivalent to the set of inflation-target policies with different speeds of adjustment. Finally, nominal-income targets are highly inefficient: they create great volatility in both inflation and output.

I. Introduction

According to mainstream economics, the main long-run role of monetary policy is to determine the level of inflation. Policy does not affect output in the long run. In the short run, however, policy influences the variability of

*This paper is a revised version of NBER Working Paper #5952. I am grateful for comments and assistance from Bob Buckle, David Mayes, Adam Posen, David Romer, the referees, numerous seminar participants and especially Hailong Qian. Financial support was provided by the Reserve Banks of Australia and New Zealand.

both output and inflation around their average levels. A good monetary policy is one that produces low average inflation, and also keeps output and inflation as stable as possible.

How can these goals be accomplished? Many economists argue that they are best achieved if central banks follow a simple policy rule. There is no consensus, however, about what rule would be best. In the 1960s and 1970s, there was strong interest in a rule prescribing a fixed growth rate for the money supply. Interest in money-growth rules has waned, however, because of apparent instability in money demand. Historically, many countries have based policy on exchange-rate targets, but few economists currently advocate this approach for major developed countries. Consequently, the last 20 years have seen a search for new policy rules among both academic economists and central bankers.

This search has produced several leading candidates. Within central banks, there is a strong movement toward targets for the level of inflation. New Zealand and Canada adopted inflation targets in the early 1990s, and seven other countries have followed (Kahn and Parrish 1998). While some academic researchers support this approach, others advocate a policy that has not yet been tried: targets for the growth of nominal GDP (for example McCallum 1993, Hall and Mankiw 1994). A final group advocates a 'Taylor rule' in which interest rates are adjusted in response to movements in output and inflation (such as Krugman 1996). No central bank has explicitly adopted a Taylor rule, but the behaviour of interest rates in the United States and several other countries closely mimics their behaviour under such a rule (Taylor 1993). In these countries, central banks raise the real interest rate by 50 basis points for each 1% rise in output or 1% rise in inflation.

Which of these policy rules is best? Any of them guarantees a low average level of inflation if the parameters of the rule are set appropriately (for example, a sufficiently low target for nominal GDP growth). The differences between the policies concern their implications for short-run fluctuations in output and inflation. This paper asks which rule does best at stabilizing these variables.

I address this question in a simple macroeconomic model consisting of two familiar equations. The first equation is a dynamic IS relation: output depends on lagged output, a lagged interest rate and a demand shock. The second is an accelerationist Phillips curve: the change in inflation depends on lagged output and an inflation shock. The model is similar to models in many undergraduate textbooks, except for the explicit inclusion of time lags (which are important for the results). The model is also similar in spirit to the more complicated macroeconometric models developed by many central banks.

In the model, the goals of stabilizing output and stabilizing inflation are not entirely compatible. As emphasized by Taylor (1994), the choice of a policy rule involves a trade-off between the variances of these two variables. Given this trade-off, I consider a policy-maker who minimizes a weighted sum of the variances; the weights are determined by his views about the relative costs of output and inflation fluctuations. The solution to this problem is the 'optimal' rule for the policy-maker.

Different policy-makers put different weights on output and inflation variances, and so it is difficult to single out one policy as optimal. I therefore focus on the weaker criterion of 'efficiency'. An efficient policy is one that minimizes the sum of output and inflation variances for some choice of weights. Equivalently, an efficient policy is one that puts the economy on the inflation variance/output variance frontier. If a policy is not efficient, then there is an alternative policy that reduces both variances, and thus is preferred by any policy-maker. It turns out that efficiency is a sufficiently strong criterion to rule out many policies in favour of a fairly narrow range of options.

The model produces three sets of results about the efficiency of alternative policy rules. First, the model produces a mixed verdict about Taylor rules. Any efficient policy can be expressed as such a rule – as a rule for setting the interest rate as a function of output and inflation. Thus the model provides a justification for this general kind of policy. However, the model yields negative results about the particular Taylor rules used in practice. For a realistic calibration of the model, the usual Taylor-rule coefficients of one half for both output and inflation are *not* efficient. Instead, efficiency requires a larger coefficient on output. According to the model, a larger output coefficient would reduce the variances of both output and inflation by dampening inflationary booms and disinflationary recessions.

The second set of results concerns inflation targeting. The properties of this policy depend on exactly how it is defined. A narrow definition is that policy tries to hit a fixed inflation target as closely as possible. When shocks push inflation away from the target, policy brings it back as quickly as possible. In the model, this policy produces a large output variance. It is efficient because it minimizes the variance of inflation, putting the economy at a corner of the inflation variance/output variance frontier. But the policy is optimal only in the extreme case that policy-makers put no weight on output.

Inflation targeting has more general appeal if it is defined more broadly. Under the broader definition, policy brings inflation back to target after a shock, but not necessarily at the maximum speed. Such a gradual-adjustment policy appears closer than the narrow policy to the actual practice of inflation targeters (Mishkin and Posen 1997). A broadly defined inflation target is optimal for any weights on output and inflation variances. As one varies the weights, the optimal policies trace out inflation targets with adjustment speeds from zero to the maximum. Thus the set of efficient policies, which can be defined by certain Taylor rules, can also be defined as the set of inflation targets with different degrees of gradualism. In the model, inflation targets and efficient Taylor rules are equivalent policies.

The final set of results concerns nominal-income targeting. In my model, this policy is highly inefficient: it creates great volatility in both output and inflation. The crucial flaw in income targeting is that it delays the return of output to trend after it is pushed away. For example, if output is below trend and inflation is near its long-run level, a recovery cannot occur, because rapid real growth would cause nominal-income growth to exceed its target. This problem leads to unappealing dynamics. When policy-makers target nominal income, shocks to the economy set off oscillations in output and inflation, in which these variables repeatedly overshoot their long-run levels.

The analysis in this paper is closely related to Svensson (1997). Indeed, the essential assumptions of my model and Svensson's are the same. The papers differ primarily because they address different questions. For example, Svensson considers issues such as monitoring adherence to policy rules, and the difference between inflation targeting and money targeting; he does not examine the efficiency of current Taylor rules or the properties of nominal-income targets. The two papers overlap in their discussions of inflation targets. In particular, Svensson also shows the optimality of gradual adjustment toward a target. Section II of this paper presents the model, Sections III–V present the three sets of results, and Section VI concludes.

II. The Model

A. Assumptions

The economy is described by two equations. The first is a dynamic IS or aggregate-spending equation:

$$y = -\beta r_{-1} + \lambda y_{-1} + \varepsilon, \ \beta > 0, 0 \pm \lambda \pm 1 , \tag{1}$$

where y is the gap between output and potential output, r is the deviation of the real interest rate from its equilibrium level, and ε is a white-noise shock. According to (1), a rise in the interest rate reduces output with a one-period lag. Output also depends on lagged output – there is 'persistence'. The shock ε captures other influences on spending, such as consumer confidence and fiscal policy.

The second equation is an accelerationist Phillips curve:

$$\pi = \pi_{-1} + \alpha y_{-1} + \eta, \ \alpha > 0 , \tag{2}$$

where π is the deviation of inflation from its average level and η is another white-noise shock. According to (2), the change in inflation depends on the output gap in the previous period. The shock η is an inflation or 'supply' shock, arising for example from large changes in commodity prices. Note that, in defining π as the deviation from average inflation, I take average inflation as given. This inflation rate is chosen by policy-makers, but I assume here that this choice has already been made (based on considerations outside the model).

I assume that the monetary authority uses the interest rate as its policy instrument. Thus the interest rate in the model is most realistically interpreted as an overnight rate. Policy-makers set the interest rate after observing the current shocks ε and η. Note that I assume that policy-makers set the *real* interest rate. In practice, the interest rates controlled directly by policy-makers are nominal rates. However, policy-makers can move the real interest rate to their desired level by setting the nominal rate equal to the desired real rate plus inflation.[1]

The model captures some of the main problems facing policy-makers in practice. Both output and inflation are subject to unforecastable shocks. Inflation is inertial: barring a favourable supply shock, policy-makers can reduce inflation only by reducing output. Finally, policy influences the economy with lags. It takes a period for policy to affect output, and a period for output to affect inflation; combining these lags, it takes two periods for policy to affect inflation. This structure captures the stylized fact that policy affects output more quickly than it affects inflation (McCallum 1995).

B. Calibration

In studying the model, I will interpret a time period as a year. With this interpretation, the lags in the model are roughly realistic. Empirical studies suggest that the major effects of interest-rate shifts on output occur in about a year, and that the major effects on inflation occur in about two years (Christiano, Eichenbaum and Evans 1996).

In parts of the analysis, it will be useful to have baseline values for the model's parameters. I choose these values based on empirical evidence for the USA. To choose α, the Phillips-curve slope, I use the fact that the sacrifice ratio for reducing inflation is $1/\alpha$. A plausible sacrifice ratio is 2.5, implying $\alpha = 0.4$ (Ball 1994a). In the aggregate-spending equation, the coefficient λ on lagged

[1]It is, of course, a simplification to assume that spending depends directly on the overnight interest rate (equation (1)). Future work might consider an extension in which spending depends on a longer-term interest rate, which is linked to the overnight rate through the term structure (as in Fuhrer and Moore 1995).

output should be large, as output fluctuations are highly persistent. I assume $\lambda = 0.8$, which is close to the lagged-output coefficient when output is regressed on measures of monetary policy (DeLong and Summers 1988). Finally, I assume that the parameter β is 1.0: a one-point rise in the interest rate reduces output by 1%. This is close to the effect in both macroeconometric models and vector autoregressions (Rudebusch 1995). With $\beta = 1.0$, a three-point rise in the interest rate reduces annual growth from 2.5% to –0.5%.

C. *Why This Model?*

In the last decade, a large literature has considered alternative rules for monetary policy. Most of this literature examines medium- or large-scale macroeconomic models that are analysed numerically; for example, the papers in Bryant, Hooper and Mann (1993) and Taylor (1999). In contrast, I consider a model with only two simple equations. My approach is meant to be complementary to work with larger models. A simple model yields sharp results with clear economic explanations. Larger models are more accurate quantitatively, but harder to understand. One can derive intuitive results with my model, and then use more detailed models to check the robustness of the results.

A related modelling issue is the treatment of expectations. My model is 'backward-looking', with current variables depending only on lagged variables and current shocks. Here I depart from recent models that include expectations of future output and inflation in equations (1) and (2); for example McCallum and Nelson (1999); Rotemberg and Woodford (1997). These forward-looking models have strong theoretical foundations, but they fail to fit key facts. In particular, they do not produce the inertia in inflation that appears in the data (Estrella and Fuhrer 1998). Models that lack inflation inertia produce misleading prescriptions for policy. For example, in a forward-looking version of my model, Zarutskie (1999) shows that the optimal policy involves *no* response of real interest rates to inflation. A rise in inflation does not require tighter policy because, without inertia, inflation disappears by itself. In my view, this kind of result should not be taken seriously by policymakers.

The ideal model would start with the microfoundations of forward-looking models and also capture inflation inertia. Researchers are searching for such a model, but there is not yet a consensus on what it is. For now, we must choose among imperfect models based on the goals of our research. Since this paper focuses on practical policy questions, I choose a model that emphasizes realism over microfoundations.

III. Taylor Rules

A. *Efficient Policies*

The optimal policy rule is the one that minimizes the loss function

$$Var(y) + \mu Var(\pi), \tag{3}$$

where *Var* denotes a variance and μ is a parameter giving policy-makers' distaste for inflation variance relative to output variance. This section derives the optimal policy for a given μ. Varying μ from zero to infinity defines the set of efficient policies.

One way to find the optimal policy is to posit a general rule, derive the variances of output and inflation as functions of the rule's parameters, and then find the parameters that minimize (3). However, it is easier to first derive some restrictions on policy that are required for efficiency. I therefore proceed as follows.

A policy rule is a rule for setting the interest rate r. The interest rate affects the economy because it determines expected output in the following period. Specifically, equation (1) implies that expected output is:

$$E[y_{+1}] = -\beta r + \lambda y \tag{4}$$

where I use the assumption that ε is unforecastable. The policy-maker can adjust Ey_{+1} to any level he chooses by setting the appropriate r given (4) and the current value of y. We can therefore view the policy-maker's problem as choosing Ey_{+1}. Given a rule for Ey_{+1}, equation (4) determines the implied rule for r.

When the policy-maker chooses Ey_{+1}, what variables are relevant? Inspection of (1) and (2) shows there is only one state variable. This variable is expected inflation in the next period, $E\pi_{+1}$, which equals $\pi + \alpha y$. When the policy-maker chooses Ey_{+1}, he takes $E\pi_{+1}$ as given because it affects inflation only after two periods. The future of the economy is determined by the state variable $E\pi_{+1}$, the policy-maker's rule for Ey_{+1}, and future shocks. Thus policy-makers set Ey_{+1} as a function of $E\pi_{+1}$. Since the model is linear-quadratic, one can show that the optimal policy is linear. Thus:

$$E[y_{+1}] = -qE[\pi_{+1}] = -q(\pi + \alpha y) \tag{5}$$

where q is a parameter to be determined (which will be positive). Combining this equation with (4) yields the interest-rate rule:

$$r = [(\lambda + \alpha q)/\beta]y + [q/\beta]\pi. \tag{6}$$

This equation is a Taylor rule: the interest rate responds positively to output and inflation.

It remains to determine the parameter q. The Appendix derives the variances of output and inflation when the interest rate is set according to (6), and finds the q that minimizes the loss function (3). The result is:

$$q = [-\mu\alpha + \sqrt{\mu^2\alpha^2 + 4\mu})]/2 \tag{7}$$

According to (7), $q = 0$ if $\mu = 0$ – if the policy-maker puts zero weight on inflation variance. q rises as μ rises, and q approaches $1/\alpha$ as $\mu \rightarrow \infty$ (by L'Hospital's rule). Thus the set of efficient policies in the model is given by the interest-rate rule (6) with q ranging from zero to $1/\alpha$.[2]

B. Interpretation

Recent research shows that Taylor rules provide a good fit to actual policy in several countries since the mid-1980s. These countries include the USA (Taylor 1993, 1994), the UK (Stuart 1996), and Germany and Japan (Davies, Brookes and Culverhouse 1996). Observers such as Krugman (1996) praise Taylor rules for 'leaning against the wind' in a sensible way: when either output or inflation moves away from its long-run level, policy-makers adjust interest rates to push it back. My model provides formal support for the intuition that a Taylor rule is optimal.

The model also determines the optimal coefficients in the Taylor rule. The two coefficients in equation (6) are increasing in q, which in turn is increasing in μ. Thus a larger μ – a greater distaste for inflation variance – means that policy-makers adjust interest rates more aggressively in response to both output and inflation. Note that policy-makers react to output movements even if $\mu \rightarrow \infty$, so they care only about inflation variance (indeed, this is the case when they react most strongly). Even if policy-makers do not care about output for its own sake, they react to output movements because output affects next period's inflation through the Phillips curve.

While the Taylor-rule coefficients vary with policy-makers' tastes, efficiency places restrictions on what they can be. Given the parameters of the model – α, β and λ – there are two kinds of restrictions. First, the coefficients must lie in certain ranges. Since the optimal q lies in $(0, 1/\alpha)$, equation (5) implies that the output coefficient lies in $(\lambda/\beta,(\lambda + 1)/\beta)$, and the inflation coefficient lies

[2]Using dynamic programming, Svensson (1997) derives the interest-rate rule that minimizes the present value of the loss function (3). As the discount factor approaches one, Svensson's solution approaches the solution given by (6) and (7).

in $(0, 1/\alpha\beta)$. Second, since q determines both coefficients, there is a restriction across the coefficients. If the inflation coefficient is k_π, the output coefficient must be $\alpha k_\pi + \lambda/\beta$.

Taylor (1994) also derives the optimality of a Taylor rule. His model, however, produces very different coefficients for the rule: both coefficients approach infinity, regardless of policy-makers' tastes. This implausible result arises because Taylor assumes there is no lag in the aggregate-spending equation: policy-makers control *current* output by adjusting the interest rate. Under this assumption, extremely aggressive policy eliminates completely the effects of demand shocks (Ball 1994b). My model produces more realistic results because it captures the lags in the effects of policy.

C. Is Current Policy Efficient?

Empirical work on Taylor rules has uncovered a striking regularity: in the countries where Taylor rules fit policy, the coefficients on both output and inflation are about one half. The past literature says little, however, about whether these coefficients are optimal, or how one could decide. (Aside from aesthetic appeal, there is no obvious reason why the coefficients should be equal or should sum to 1.) If one calibrates the model of this paper, it is possible to compare actual and efficient policies.

Recall the base parameter values for the model: $\alpha = 0.4$, $\beta = 1.0$, and $\lambda = 0.8$. For these values, efficiency requires an inflation coefficient in the range (0, 2.5) and an output coefficient in (0.8, 1.8). The actual inflation coefficient of 0.5 is in the efficient range, but the output coefficient is too low. In addition, for an inflation coefficient of 0.5, the cross-coefficient restriction implies an output coefficient of 1.0, twice the actual value. Thus current Taylor rules are inefficient. Efficiency requires that policy-makers respond more aggressively to output movements than they actually do.

Why is a low output coefficient inefficient? Recall that a rule for choosing r is equivalent to a rule for choosing expected output in the next period. In this rule, the coefficient on current output is $-\beta k_y + \lambda$, where k_y is the output coefficient in the interest-rate rule. For $k_y = 1/2$ and base values of β and λ, this expression is positive. Thus a positive output gap implies a positive expected gap in the next period: the interest rate does not rise enough to eliminate output persistence. This fact is the source of inefficiency. Persistence leads to larger output fluctuations, which generate larger inflation fluctuations through the Phillips curve. Policy-makers should raise the output coefficient in the Taylor rule to produce a non-positive effect of y on Ey_{+1}.

This conclusion is reasonably robust to variation in parameter values. For example, there is uncertainty about the parameter β, the effect of the interest rate on output, and β may vary across countries. However, for the base values of α and λ, an output coefficient of 0.5 is below the efficient range for any $\beta < 1.6$. The coefficient is below the level implied by the cross-coefficient restriction for any $\beta < 2.67$. This is a generous bound: $\beta = 2.67$ implies that a 1-point rise in the interest rate pushes growth from its trend level to below zero.

The conclusion that countries with Taylor rules should respond more aggressively to output fluctuations is strengthened by examining recent history. Consider the UK, where a Taylor rule with one-half coefficients fits policy since 1986. From 1986 through 1990, the UK experienced large positive output gaps, and CPI inflation rose from 3.4% to 9.5% (OECD 1996). Then the gaps turned negative through 1995, and inflation ended up back at 3.4%. This boom-bust cycle appears inefficient. If the UK's Taylor rule had a larger coefficient on output, policy would have been tighter in the late 1980s and looser in the 1990s. This would have dampened the output fluctuations and the temporary run-up in inflation. The same story fits the USA, although the US boom-bust cycle was smaller.[3]

IV. Inflation Targets

This section considers policies that attempt to achieve a target level of inflation. This goal implicitly defines a rule for setting the interest rate. Inflation targets are of interest because a number of countries have adopted them during the 1990s. What relation do these policies have to the efficient policies derived in the previous section?

A. Strict Inflation Targets

A strict definition of inflation targeting is that policy minimizes the variance of inflation around its average level. As discussed by Svensson (1997), this implies that policy sets to zero the expected inflation deviation in two periods:

$$E\pi_{+2} = 0. \tag{8}$$

[3]The more desirable outcomes for the USA reflect the fact that, while policy generally followed a Taylor rule with one-half coefficients, it deviated at two important points. In 1986, interest rates were much higher than those implied by the rule; in 1992–93, they were lower. Thus the USA leaned more heavily against the 1980s boom and 1990s recession than the Taylor rule dictated. Some observers interpret these deviations as mistakes or responses to special circumstances (*The Economist* 1996). My model suggests that actual policy was preferable to the standard Taylor rule.

Recall that policy cannot affect inflation before two periods because of the lags in the model.

A strict inflation target is an efficient policy by definition. It minimizes the objective function (3) in the limiting case of $\mu \rightarrow \infty$ (policy-makers care only about inflation variance). The interest-rate rule implied by strict inflation targeting is the Taylor rule (6) with q at its upper bound of $1/\alpha$. For my base parameter values, this implies an output coefficient of 1.8 and an inflation coefficient of 2.5.

While a strict inflation target is efficient, it is optimal only for extreme tastes. If policy-makers care about both output and inflation, they should move interest rates less aggressively. In general, it is not clear whether a strict inflation target would improve on current Taylor rules with coefficients of one half. For base parameter values, one can show that the inflation target would reduce the variance of inflation but raise the variance of output. A policy-maker prefers the status quo if the taste parameter μ is sufficiently small.

B. *Inflation Targets with Gradual Adjustment*

Inflation targeting can be made more attractive by defining it more broadly. We can see this by rewriting the condition for an efficient rule from Section III. Note first that the Phillips curve (2) implies an expression for expected inflation in two periods:

$$E\pi_{+2} = E\pi_{+1} + \alpha E y_{+1}. \tag{9}$$

According to (5), efficiency requires that Ey_{+1} be set equal to $-qE\pi_{+1}$. Substituting this condition into (9) yields:

$$E\pi_{+2} = (1 - \alpha q)E\pi_{+1}. \tag{10}$$

One can define the set of efficient policies by equation (10) with q ranging from 0 to $1/\alpha$.

Equation (10) is a generalization of a strict inflation target, equation (8). Instead of setting $E\pi_{+2}$ equal to 0, the rule sets it equal to a fraction of $E\pi_{+1}$. Given the range for q, this fraction can range from 0 to 1. Thus equation (10) is a partial-adjustment variation on an inflation target. Given expected inflation in one period, which policy-makers take as given, policy moves expected inflation part of the way to its long-run target in the following period.

It makes sense that a partial-adjustment rule is the optimal policy. Policy-makers want to return inflation to its target, but they face a quadratic cost of changing inflation: a change in inflation requires a deviation of output from potential, and the square of this deviation enters the loss function. Recall that policy-makers who put a large weight on output variance choose a small value of the parameter q. In equation (10), this means they adjust inflation slowly toward its target. Slow adjustment raises the variance of inflation but reduces the variance of output.

If we define inflation targeting to allow partial adjustment, then we have a strong result: the class of efficient policies is equivalent to the class of inflation targets. The efficient Taylor rules derived above are the interest-rate rules that implement inflation targets with various adjustment speeds. Since any efficient rule can be interpreted as inflation targeting, any policy-maker should choose this policy. If he cares a lot about output, he should choose a slow adjustment speed, but he can still define policy entirely by a path for inflation. When policy is expressed as an inflation target, there is no need to mention output or other variables.

An inflation target with gradual adjustment is probably a good description of policy in some countries. For example, Australia has adopted an inflation target of 2–3% to be met 'on average over the cycle' rather than continuously (Reserve Bank of Australia 1996). This rule appears to allow gradual adjustment after a shock. As discussed by Mishkin and Posen (1997), few if any 'inflation target' countries have strict targets as defined here. In choosing a short-run path for inflation, policy-makers put some weight on the consequences for output.

V. Nominal Income Targets

Many economists advocate targets for nominal income. Authors such as Hall and Mankiw (1994) and McCallum (1993, 1995) argue that this policy would produce better outcomes, and particularly more stable output, than inflation targets. So far, however, no central bank has chosen to target nominal income. My model suggests that this is wise: nominal-income targets produce great *instability* in output and inflation.

The literature discusses two kinds of nominal-income targets, growth targets and level targets. Growth targets mean that policy always aims for the same rate of income growth, regardless of past growth. Level targets mean that policy reverses past mistakes to return the income level to a fixed path. I focus here on income-growth targets, which have received the most attention. An earlier version of this paper (Ball 1997) considers level targets, and shows that they have equally poor properties.

A. The Overshooting Problem

An income-growth policy minimizes the variance of income growth. Equivalently, it sets expected income growth in the next period equal to a fixed target. Note that policy can influence income growth in one period rather than two, because it affects output with a one-period lag. In deviations from trend, income growth is the sum of output growth, $y_{+1} - y$, and inflation, π_{+1}. Thus an income-growth policy is defined by:

$$E[y_{+1} - y + \pi_{+1}] = 0 \tag{11}$$

To determine the interest-rate rule under income targeting, I substitute (1) and (2) into (11). This leads to:

$$r = [(\alpha + \lambda - 1)/\beta]y + [1/\beta]\pi. \tag{12}$$

Substituting this rule into (1) yields:

$$y = (1 - \alpha)y_{-1} - \pi_{-1} + \varepsilon \tag{13}$$

This equation and the Phillips curve (2) define the process for output and inflation under income targeting.

To understand the behaviour of output and inflation, it helps to examine the impulse-response functions implied by (2) and (13). Specifically, consider the effects of a unit shock to the aggregate-spending equation: $\varepsilon = 1$ in period 0. In this experiment, I set the parameter α equal to 1, which yields the simplest dynamics (as shown below, the qualitative results generalize).

Figure 1 shows the effects of the spending shock. By construction of the policy, nominal-income growth is unaffected after the period when the shock occurs. But output and inflation become highly unstable. After initially rising, each falls, overshooting its long-run level, and then oscillates indefinitely.

By contrast, Figure 2 shows the responses to a spending shock under strict inflation targeting. In this case, the effects of the shock are eliminated quickly. Clearly income targeting produces much greater variability in both output and inflation in response to aggregate-spending shocks. Similar calculations show that the same is true for shocks to the inflation equation.

Why does income targeting produce these strange dynamics? The output shock in period 0 causes inflation to rise in period 1. Policy then causes a recession, bringing inflation back to its long-run level in period 3. This is the point at which the crucial flaw in income targeting arises. With the inflationary effect of the initial shock eliminated, it would be optimal to let output recover. But recovery to trend would require high output *growth*; with inflation on target, this would imply above-target growth in nominal income. Thus a

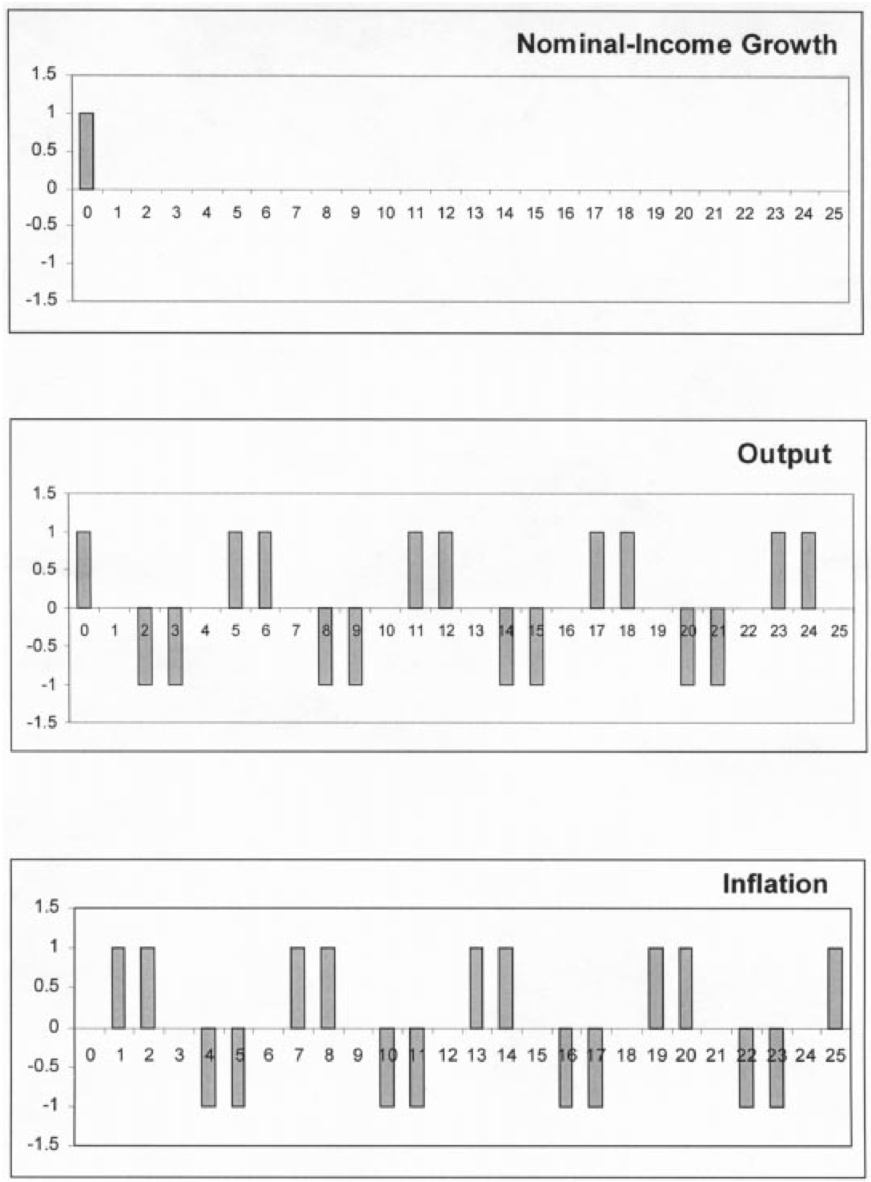

Figure 1: Nominal-income targeting: responses to an output shock

nominal-income rule prevents output from recovering: it requires that the level of output stays low while inflation is at its target. Because the recession continues in period 3, inflation falls *below* target in period 4. Policy then creates a boom to raise inflation, but it overshoots the target again for analogous reasons. The oscillations in Figure 1 are underway.

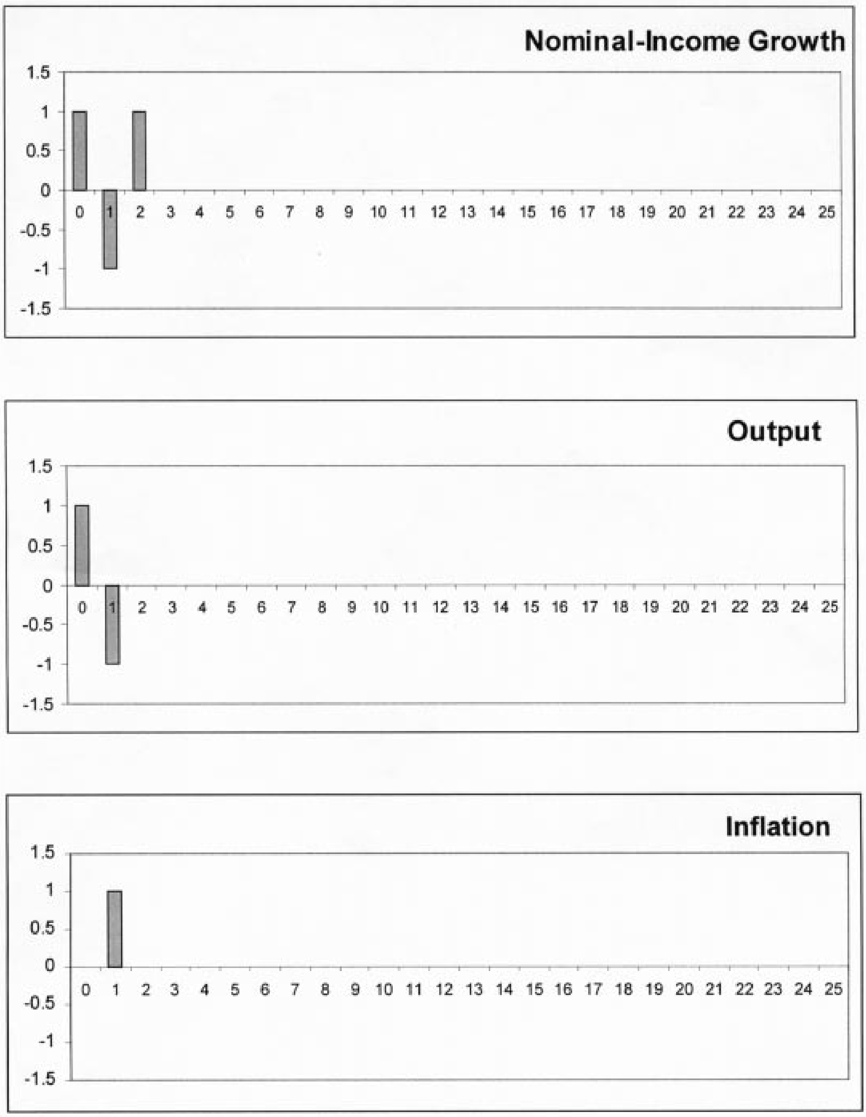

Figure 2: Strict inflation targeting: responses to an output shock

B. Infinite Variances?

So far I have used an example to illustrate the instability caused by nominal-income targeting. Part B of the Appendix derives the general properties of income targeting, giving a surprising result: under income targeting, the variances of both output and inflation are infinite. Thus, in my model, income targeting is not just inefficient, but disastrous. It causes both output and inflation to wander arbitrarily far from their long-run levels. This result holds for any values of the model's parameters.

As shown in the Appendix, this result arises because income targeting causes inflation and output to follow non-stationary processes. One can see intuitively that these variables are non-stationary from the impulse responses in Figure 1. In the figure, the oscillations in output and inflation never die out: the effects of a single shock last forever. This fact implies non-stationarity, which in turn implies infinite variances.

The result that output and inflation have infinite variances may appear too extreme to believe. Indeed, recent work by McCallum (1997) and Dennis (1998) shows that the result is not robust. Reasonable modifications of the Phillips curve, equation (2), produce finite variances of output and inflation under income targeting. Dennis, for example, shows that the variances are finite if inflation depends on both past inflation and expected future inflation. This is true even if the future-inflation term has a small weight, so the Phillips curve is close to my purely backward-looking equation.

These results do not, however, overturn the conclusion that income targeting is an undesirable policy. While infinite variances are a special result, it appears that income targeting produces finite but large variances under a broad range of assumptions. The reason is the overshooting phenomenon discussed above. The overshooting result appears robust; it arises, for example, in Dennis's model, as long as the lagged-inflation term in the Phillips curve has a positive weight. In that model, the oscillations die out over time, implying finite variances. But the variances are still large relative to efficient policies.[4]

A key policy question is the relative performance of nominal-income targeting and inflation targeting. I conjecture that inflation targeting produces smaller variances for both output and inflation in a wide range of models, because it does not cause overshooting. Future work should examine this conjecture in generalizations of my model.

VI. Conclusion

The model in this paper yields several prescriptions for monetary policy. Countries such as the USA are right to follow Taylor rules, but they should change the coefficients to respond more strongly to output fluctuations.

[4]Taylor's (1985) simulations of a nominal-income rule also produce oscillations that die out over time. In Taylor's model, the Phillips curve includes only past inflation, but the coefficient on this term is less than one.

Alternatively, countries should adopt inflation targets with gradual adjustment; such a policy is equivalent to an efficient Taylor rule. Finally, policy-makers should *not* adopt nominal-income targets.

The model in this paper has recently been extended in a number of directions. Rudebusch and Svennson (1999) study a model that is similar in spirit to both this paper and Svensson (1997), but more detailed. Their model is quarterly, and the parameters are estimated from US time-series data. Rudebusch and Svennson largely confirm the results of this paper. They too find that the Taylor-rule coefficient on output should be raised above one half.

Ball (1999) extends the model to an open economy. In the extended model, monetary policy affects the economy through both interest-rate and exchange-rate channels. This changes the optimal policy rule in two ways. First, the instrument on the left side of the Taylor rule becomes a 'Monetary Conditions Index': an average of the real interest rate and the real exchange rate. Second, on the right side of the rule, inflation is replaced by 'long-run' inflation, a variable that filters out the transitory effects of exchange-rate fluctuations. If policy is expressed as a target rule, the target variable is long-run inflation.

This paper and the related literature raise questions for future research. One puzzle is why central banks follow Taylor rules with coefficients of one half when current models suggest that larger coefficients are optimal. One possibility is that one-half coefficients are optimal in reality for reasons left out of the models. One phenomenon that most models ignore is uncertainty about the structure of the economy and the values of parameters. Some authors argue that uncertainty makes optimal policy less aggressive, which suggests smaller Taylor-rule coefficients (for example, Blinder 1997). It is not yet clear whether models with uncertainty can rationalize the coefficients observed in practice.[5] Alternatively, current policy may truly be suboptimal. There could be political pressures that lead policy-makers to adjust interest rates more cautiously than they should. For example, Lowe and Ellis (1997) suggest that central bankers are embarrassed by reversals in interest-rate movements, because they look like repudiations of earlier actions. Policy-makers may move interest rates less than they think is optimal to assure that future adjustments are in the same direction.

Another issue requiring further research is the optimal speed of adjustment toward an inflation target. In the model of this paper, this adjustment speed (or equivalently, the parameter q in the Taylor rule) is the one aspect of policy that is not determined by efficiency alone. Instead, it is determined by policy-makers' preferences. Policy-makers adjust quickly if they care mainly

[5] Recent work on this issue includes Sargent (1999) and Rudebusch (1998).

about inflation variability, and slowly if they care mainly about output variability.

It is not clear, however, whether this result is robust. The choice of adjustment speed may have effects that my model ignores; in particular, it may influence the output–inflation trade-off. The model assumes that the sacrifice ratio for reducing inflation is fixed at $1/\alpha$. But the empirical results in Ball (1994a) suggest that the sacrifice ratio is smaller if disinflation is quick. This effect increases the optimal speed of adjustment to an inflation target. If the effect is large enough, then output variance is minimized by immediate adjustment – a strict inflation target – and this policy dominates all others. More research on the Phillips curve is needed to determine the importance of this point.

Laurence Ball
Department of Economics
Johns Hopkins University
Baltimore, MD 21218
USA

References

Ball, Laurence (1994a), 'What Determines the Sacrifice Ratio?', in N. Gregory Mankiw, ed., *Monetary Policy*. Chicago: University of Chicago Press.

—— (1994b), 'Discussion', in Jeffrey C. Fuhrer, ed., *Goals, Guidelines and Constraints Facing Monetary Policymakers*. Boston: Federal Reserve Bank of Boston.

—— (1997), 'Efficient Rules for Monetary Policy', NBER Working Paper No. 5952.

—— (1999), 'Policy Rules for Open Economies', in John B. Taylor ed. *Montary Policy Rules*. Chicago: University of Chicago Press.

Blinder, Alan S.(1997), *Central Banking in Theory and Practice*. Cambridge, MA: MIT Press.

Bryant, Ralph, Peter Hooper and Catherine Mann (eds) (1993), *Evaluating Policy Regimes: New Research in Empirical Macroeconomics*. Washington DC: The Brookings Institution.

Christiano, Lawrence J., Martin Eichenbaum and Charles Evans (1996), 'The Effects of Monetary Policy Shocks: Evidence from the Flow of Funds', *Review of Economics and Statistics* 78(February), 16–34.

Davies, Gavyn, Martin Brookes and Stuart Culverhouse (1996), Goldman Sachs report, London: Goldman Sachs.

DeLong, J. Bradford and Lawrence H. Summers (1988), 'How Does Macroeconomic Policy Affect Output?', *Brookings Papers on Economic Activity*, 2, 433–80.

Dennis, Richard (1998), *Instability Under Nominal GDP Targeting: The Role of Expectations*. Unpublished paper. Canberra: Australian National University.

Estrella, Arturo and Jeffrey C. Fuhrer (1998), *Dynamic Inconsistencies: Counterfactual Implications of a Class of Rational Expectations Models*. Unpublished paper. Boston: Federal Reserve Bank of Boston.

Fuhrer, Jeffrey C. and George R. Moore (1995), 'Monetary Policy Trade-offs and the Correlation Between Nominal Interest Rates and Real Output', *American Economic Review,* 85(March), 219–39.

Hall, Robert E. and N. Gregory Mankiw (1994), 'Nominal Income Targeting', in N. Gregory Mankiw, ed., *Monetary Policy*. Chicago: University of Chicago Press.

Hendry, David F. (1995), *Dynamic Econometrics*. Oxford: Oxford University Press.

Kahn, George A. and Klara Parrish (1998), 'Conducting Monetary Policy with Inflation Targets', Federal Reserve Bank of Kansas City *Economic Review*, 83(3rd Quarter), 5–32.

Krugman, Paul (1996), 'Stable Prices and Fast Growth: Just Say No', *The Economist*, August 31, 15–18.

Lowe, Philip and Luci Ellis (1997), 'The Smoothing of Official Interest Rates', in Philip Lowe, ed., *Monetary Policy and Inflation Targeting*. Sydney: Reserve Bank of Australia.

McCallum, Bennett T. (1993), 'Specification and Analysis of a Monetary Policy Rule for Japan', Bank of Japan *Monetary and Economic Studies,* 11, 1–45.

—— (1995), 'New Zealand's Monetary Policy Arrangements: Some Critical Issues', Reserve Bank of New Zealand Discussion Paper G95/4, June.

—— (1997), 'The Alleged Instability of Nominal Income Targeting', NBER Working Paper No. 6291.

—— and Edward Nelson (1999), 'Performance of Operational Policy Rules in an Estimated Semi-Classical Structural Model', in John B. Taylor ed. *Monetary Policy Rules*. Chicago: University of Chicago Press.

Mishkin, Frederic S. and Adam S. Posen (1997), 'Inflation Targeting: Lessons from Four Countries', Federal Reserve Bank of New York *Economic Policy Review*, August.

Organization for Economic Cooperation and Development (1996), *Economic Outlook*, 59(June).

Reserve Bank of Australia (1996), 'Statement on the Conduct of Monetary Policy', *Bulletin*, September.

Rotemberg, Julio J. and Michael Woodford (1997), 'An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy', *NBER Macroeconomics Annual*, 297–361.

Rudebusch, Glenn (1995), 'What Are the Lags in Monetary Policy?', *FRBSF Weekly Letter*, (February 3), 95–105.

—— (1998), *Is the Fed Too Timid? Monetary Policy in an Uncertain World*. Unpublished paper. San Francisco: Federal Reserve Bank of San Francisco.

—— and Lars E. O. Svensson (1999), 'Policy Rules for Inflation Targeting', in Taylor,

Sargent, Thomas (1999), 'Discussion', in John B. Taylor ed. *Monetary Policy Rules.* Chicago: University of Chicago Press.

Stuart, Alison (1996), 'Simple Monetary Policy Rules', *Bank of England Quarterly Bulletin* 36 (August), 281–7.

Svensson, Lars E. O. (1997), 'Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets', *European Economic Review*, 41, 1111–46.

Taylor, John B. (1985), 'What Would Nominal GNP Targetting Do to the Business Cycle?', *Carnegie-Rochester Conference Series on Public Policy*, 22(Spring), 61–84.

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

—— (1994), 'The Inflation / Output Variability Trade-off Revisited', in Jeffrey C. Fuhrer, ed., *Goals, Guidelines, and Constraints Facing Monetary Policymakers*. Boston: Federal Reserve Bank of Boston.

—— (ed.) (1999), *Monetary Policy Rules*. Chicago: University of Chicago Press.

The Economist (1996), 'Economics Focus', August 10, 60.

Zarutskie, Rebecca (1999), *Efficient Monetary Policy in a Forward-Looking Model.* Unpublished paper. Baltimore: Johns Hopkins University.

Appendix

A. Derivation of the Optimal Taylor Rule

Here I derive the optimal value of the parameter q in equation (6). This completes the solution for the optimal interest-rate rule.

Substituting (6) into (1) yields:

$$y = -\alpha q y_{-1} - q\pi_{-1} + \varepsilon \tag{A1}$$

This equation and (2) define a vector AR-1 process for output and inflation:

$$X = BX_{-1} + E \tag{A2}$$

where $X = [y\ \pi]'$, $E = [\varepsilon\ \eta]'$, and B is a 2×2 matrix with elements $b_{11} = -\alpha q$, $b_{12} = -q$, $b_{21} = \alpha$, and $b_{22} = 1$. The variance-covariance matrix of X, denoted V, is given in vector form by:

$$\text{vec}(V) = [I - (BXB)]^{-1}\text{vec}(\Omega) \tag{A3}$$

where Ω is the variance-covariance matrix of E (see Hendry 1995, p. 112). After taking the inverse in (A3), one can find the variances of y and π:

$$[Var(y)\ Var(\pi)]' = D[\sigma_\varepsilon^2\ \sigma_\eta^2]' \tag{A4}$$

where σ_ε^2 and σ_η^2 are the variances of ε and η and D is a 2 × 2 matrix with elements $d_{11} = 2/(2 - \alpha q)$, $d_{12} = q/(2\alpha - \alpha^2 q)$, $d_{21} = \alpha/(2q - \alpha q^2)$, and $d_{22} = (1+2\alpha q - \alpha^2 q^2)/(2\alpha q - \alpha^2 q^2)$.

The optimal value of q is the one that minimizes $Var(y) + \mu Var(\pi)$ (equation (3)). The solution to this minimization problem is given by equation (7) in the text.

B. Income Targeting

Here I show that nominal-income targeting implies infinite variances for output and inflation, as discussed in Section VB.

Under income targeting, equations (2) and (13) define a vector $AR - 1$ process for output and inflation:

$$X = CX_{-1} + E \tag{A5}$$

where C is a 2 × 2 matrix with elements $c_{11} = (1 - \alpha)$, $c_{12} = -1$, $c_{21} = \alpha$ and $c_{22} = 1$. The behaviour of y and π is determined by the eigenvalues of C:

$$z = [2 - \alpha \pm \sqrt{\alpha^{2} - 4\alpha}]/2 \tag{A6}$$

If $\alpha < 4$, the eigenvalues are complex and lie on the unit circle. This implies that output and inflation follow oscillatory processes that neither die out nor explode. If $\alpha > 4$ (a less realistic case), one eigenvalue is less than –1, and the processes are explosive. In both cases, the eigenvalues imply that y and π are non-stationary, and thus that their variances are infinite.