

This PDF is a selection from a published volume from the
National Bureau of Economic Research

Volume Title: The Inflation-Targeting Debate

Volume Author/Editor: Ben S. Bernanke and Michael Woodford,
editors

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-04471-8

Volume URL: <http://www.nber.org/books/bern04-1>

Conference Date: January 23-26, 2003

Publication Date: December 2004

Title: Inflation Targeting, Price-Path Targeting, and Output
Variability

Author: Stephen G. Cecchetti, Kim

URL: <http://www.nber.org/chapters/c9558>

Inflation Targeting, Price-Path Targeting, and Output Variability

Stephen G. Cecchetti and Junhan Kim

4.1 Introduction

The 1990s were amazing in many ways. Not only did the internet and cellular phones come into widespread use, but overall economic conditions improved nearly everywhere we look. Growth was higher, inflation was lower, and both were more stable. In the United States, for example, inflation fell from 6 percent at the beginning of the decade to less than 2 percent by the end. Meanwhile real growth rose from less than 3 percent to over 4 percent. Volatility declined, too. The American case is the most dramatic instance of what has really been a worldwide trend.¹ And while these improvements in economic performance could have been the consequence of the world being calmer, Cecchetti, Flores-Lagunes, and Krause (2002) argue that roughly three-quarters of it can be explained by better monetary policy. That is, central bankers did a better job of stabilizing inflation at low levels while keeping growth high.

Making better monetary policy is not just a problem of finding competent central bankers. In fact, there is a history of central bankers who tried to do their jobs but were thwarted by politicians. Over the years we have

Stephen G. Cecchetti is professor of international economics and finance at Brandeis University and a research associate of the National Bureau of Economic Research. Junhan Kim is an economist at the Bank of Korea.

We wish to thank our discussant Greg Mankiw, as well as Edwin Truman and Mark Wynne for providing us with data, and Bill Gavin, Stefan Krause, Lianfa Li, Roisin O'Sullivan, and the conference participants, especially Lars Svensson, for helpful suggestions. The views expressed herein are those of the authors and not necessarily those of the Bank of Korea.

1. Cecchetti and Ehrmann (2002) compare the 1985–89 period with 1993–97 for a set of twenty-three industrialized and emerging-market countries and find that annual inflation fell by an average of 5 percentage points, annual growth rose by an average of 1 percentage point, and both were significantly more stable.

learned that the institutional environment is at least as important as the people in ensuring good policy outcomes. Without a well-designed central bank, the people in charge don't have a chance. Today, we have a good sense of what best practice is in the design of central banks. First, it is crucial that monetary policymakers are independent of short-term political influences. Second, these independent central bankers must be held accountable through mechanisms that involve public announcement of objectives. Inflation targeting is the most common formulation of the sort of policy regime in place today.² The primary element of inflation targeting is a public commitment to price stability in the form of a medium-term numerical inflation target.

With the success of inflation targeting has come a discussion of potential refinements. One issue is whether the central bank should adopt a target for inflation or a target for the path of the price level. With an inflation target, the central bank simply tries to ensure that period-by-period inflation remains close to the target. When inflation turns out to be above or below the target, the miss is forgotten. Bygones are bygones, so there is a form of base drift in the (log) price level. Price-path targeting, or "price-level targeting" as it is often called, is different as it implies that when the price level is above or below the target path, the objective of policy is to return it to the present target path.³ This means that if prices move above the target path, then policy will need to bring them back down.⁴ But which one is better? Should central banks be instructed to target inflation or target the price path?

Svensson (1999) is the first person to take on this question. He starts by assuming that society cares about inflation. The social objective is to minimize the expected present discounted value of the weighted average of squared deviations of inflation and output from their targets. He then posits that the central bank can be bound to meet a particular objective but not to respond to shocks in a specific way. That is, the central bank will always have discretion in adjusting its instrument, but it can be held accountable for its objective. This sort of discretion, what we might refer to as "instrument discretion," implies that if we were to instruct central bankers to minimize the true social loss function, there would be a bias. The exact form of the bias depends on the structure of the economy, but in most cases there is a bias toward stabilizing output.⁵ One solution to this problem is to instruct the central bank to minimize a loss function that deviates

2. For a brief synopsis of what inflation targeting entails see Mishkin (forthcoming).

3. We adopt the terminology "price-path" targeting rather than the traditional "price-level" targeting to emphasize that the target path can have a positive slope and so a period of inflation need not be countered with one of deflation.

4. Mervyn King (1999) argues that in practice there is little difference between inflation targeting and price-path targeting. The reason is that politicians will hold central bankers accountable for meeting inflation targets over sufficiently long horizons so that it will look like a price-path target. We will take this up in more detail below.

5. For a discussion see Clark, Goodhart, and Huang (1999).

from society's. Rogoff (1985) suggested appointing central bankers that are more avid inflation hawks than the public at large.

In this context, Svensson shows that in countries where output is sufficiently persistent, performance can be improved by instructing policymakers to target the price path, even though society cares about inflation.⁶ To understand why output persistence is central to the result, note that the more persistent output is, the longer output stays away from equilibrium following a disturbance. Now consider the possibility of a policy response. Monetary policy responds to shocks by inducing a price-level surprise, immediately creating a conflict between the output and inflation stability objectives. And the more persistent output is, the longer-lasting the shocks and the more important it will be to respond aggressively to them. If the goal is to stabilize prices, then these aggressive responses will have to be undone quickly, which ends up lowering the volatility of inflation.

There are several issues that arise in considering this result. First, Svensson compares inflation targeting with price-path targeting in order to emphasize the contrast between the two. But there is really a continuum of intermediate possibilities that weight the two. Batini and Yates (2003) have labeled these "hybrid-targeting" regimes. We begin by showing that for a given degree of output persistence, there is an optimal hybrid-targeting policy that is a weighted average of inflation and price-path targeting. But second, and more important, the focus on output persistence means that the choice is an empirical one. What is the optimal regime for a given country? Beyond this, there is the question of whether it is worth trying to move to the optimal regime. Clarity is and should be prized in central banking.

In fact, an optimal hybrid target sacrifices simplicity for optimality. It is much more difficult to explain a hybrid than it would be to explain either of the extreme alternatives. However, as King (1999) has suggested, one of the key policy choices is the horizon over which central bankers are evaluated. That is, are they asked to maintain inflation at or near the target level on average every two, three, five, or even ten years? Put another way, central bankers will have a horizon over which they are expected to bring the price level back to its desired path. Under this interpretation, hybrid targeting becomes a statement about the *optimal horizon* over which the price level is brought back to the desired path; it may not be that hard to convince people that they should give the central bank some time to fight back unwanted price shocks.

Even so, the idea that central bankers should, for strategic reasons, be told to do something that explicitly deviates from what society truly cares about will trouble many people. Should we go to the effort of explaining that we are instructing the central bank to do one thing, while we care

6. More recent papers by Dittmar and Gavin (2000) and Vestin (2000) confirm this result.

about another, because we know that they can't be trusted? Again, this is an empirical question. How much do we lose by just telling monetary policymakers to target the thing that society cares about?

To address these issues, we examine a set of twenty-three countries and find that for nearly all of them some form of hybrid-targeting regime would be optimal—at least in principle. But we go on to show that adopting such an optimal regime has only very modest benefits (as measured by the percentage reduction in the social loss) when compared with strict inflation targeting. In other words, once you look at the numbers closely it is hard to see the benefit of starting to engage in what would surely be a very difficult public dialogue. Our conclusion is that we should hold central bankers accountable for meeting our social loss function, not some contrived one that might incrementally improve macroeconomic performance.

The remainder of this paper is organized as follows. First, we set out the theoretical problem and derive the optimal hybrid-targeting regime, and we show that this can be interpreted as the optimal horizon. We also show the relationship between output persistence and the weight on price stability. We also present a set of empirical results and compare the loss between an optimal-targeting regime and inflation or price-path targeting. The final section concludes.

4.2 Hybrid Targeting

The theoretical exercise is straightforward. Society cares about a weighted average of inflation and output deviations from their target paths. If it were possible to bind policymakers to react to shocks in a particular way, then it would be optimal to give them society's objective and then hold them accountable for adjusting their policy instrument in the way prescribed by the reaction function that minimizes this social objective. But such commitment is impossible (and may not even be desirable). Instead, the central bank can be held accountable for minimizing a loss function under discretion. What should that loss function be?

To answer this question, we proceed in two steps. First, we derive the central bank's policy reaction function, or instrument rule, under discretion for a family of loss functions that admits a wide variety of targeting regimes. Second, given the solution we find the targeting regime that minimizes the social loss. This is the optimal hybrid.

4.2.1 The Central Banker's Problem

The policymaker solves a standard optimal control problem, choosing the path of the price level that minimizes a quadratic loss function subject to the constraints imposed by the linear structure of the economy. We assume that the central bank minimizes

$$(1) \quad L^{CB} = E \left\{ \sum_t \beta^t [\lambda(p_t - p_t^*)^2 + (1 - \lambda)(y_t - y_t^*)^2] \right\},$$

where L^{CB} is the central bank's loss, E is the expectation operator, p_t is the (log) actual price level, p^* is the desired price level, y_t is the (log) actual output, y^* is desired (or potential) output level, λ is the degree to which the central bank prefers price stability to output stability, and β is the time discount factor. Equation (1) is sufficiently general to admit inflation targeting, price-path targeting, and everything in between. Targeting regimes differ depending on how the target, p_t^* , is defined. The simplest cases are inflation targeting, where

$$(2) \quad p_t^*(IT) = p_{t-1} + \pi^*,$$

and price-path targeting, where

$$(3) \quad p_t^*(PPT) = p_{t-1}^* + \pi^*.$$

In both cases the inflation target is π^* . But under inflation targeting, given by equation (2), the target is an increment over the past period's realized price level, whatever it turned out to be. By contrast, under price-path targeting, the current target is an increment over the past period's *target*.

Hybrid targeting is a weighted average of inflation and price-path targeting. That is,

$$(4) \quad \begin{aligned} p_t^*(Hybrid) &= \eta(p_{t-1} + \pi^*) + (1 - \eta)(p_{t-1}^* + \pi^*) \\ &= \eta p_{t-1} + (1 - \eta)p_{t-1}^* + \pi^*, \end{aligned}$$

where η is the weight on inflation targeting. Notice that $\eta = 1$ and $\eta = 0$ are the special cases, inflation and price-path targeting, respectively. Substituting equation (4) into the loss function (1), and normalizing various constants and initial conditions to zero, we get

$$(5) \quad L^{CB} = E \left\{ \sum_t \beta^t [\lambda(p_t - \eta p_{t-1})^2 + (1 - \lambda)y_t^2] \right\}.$$

Normalization implies that y is now the output gap and that the price path is now measured as the deviation from the inflation objective π^* .

Following Svensson (1999) and others, we assume that the dynamics of the economy are adequately described by a neoclassical Phillips Curve.⁷ That is,

7. We choose the neoclassical Phillips curve because of its theoretical tractability. There are a number of alternatives, including the now common New Keynesian Phillips curve in which the output gap depends on expected future prices rather than current ones, and the aggregate supply formulation derived by Mankiw and Reis (2001) in their work on sticky information. While it would be feasible to examine these alternatives numerically, the more conventional Phillips curve allows us to derive a wider range of conclusions.

$$(6) \quad y_t = \rho y_{t-1} + \alpha(p_t - p_t^e) + \varepsilon_t,$$

where p_t^e is the expectation of p at time t , ρ and α are constants, and ε is an independently and identically distributed (i.i.d.) shock with variance σ_ε^2 . For the points that we wish to make here, this closed-economy model is sufficient. In the empirical section, we expand the analysis to an open-economy version that includes import prices as well.

The job of the central bank is to choose a path for the price level p_t that minimizes the loss (5) subject to equation (6).⁸ Assuming rational expectations, we can use the techniques described in Svensson (1999) to first derive the first-order conditions, guess the solution, and then use the method of undetermined coefficients.⁹ The first-order conditions include the output equation (6) and

$$(7) \quad p_t - \eta p_{t-1} = -\frac{\alpha(1-\lambda)(1-\rho\beta)}{\lambda[1-\rho\beta(\rho-b\alpha)](1-\eta\rho\beta)} y_t.$$

Equation (7) embodies the trade-off between output and prices in the loss function. It tells us the extent to which prices react to output shocks along an optimal path. Under rational expectations, we know that the solution for the price level must be of the form

$$(8) \quad p_t = ap_{t-1} + by_{t-1} + c\varepsilon_t.$$

We can solve this for

$$\begin{aligned} a &= \eta \\ b &= \frac{-(1-\rho^2\beta) + \sqrt{(1-\beta\rho^2)^2 - 4\rho^2\alpha^2\beta\frac{1-\lambda}{\lambda}}}{2\rho\beta\alpha} \\ c &= -\frac{D}{1+\alpha D}, \\ \text{where } D &= \frac{\alpha(1-\lambda)}{\lambda[1-\rho\beta(\rho-b\alpha)]}. \end{aligned}$$

Setting η equal to either zero or 1, this solution collapses to the one in Svensson (1999).

This formulation allows us to write the laws of motion for output and prices, and these are

$$(9) \quad y_t = \rho y_{t-1} + (1+\alpha c)\varepsilon_t$$

8. By adding an aggregate-demand curve relating the price level to the interest rate, we could shift the problem to one in which the central bank does not choose prices directly. This increase in complexity changes none of our results.

9. See also Söderlind (1999).

$$(10) \quad p_t = \eta p_{t-1} + b y_{t-1} + c \varepsilon_t.$$

That is, output depends on lagged output, while prices depend on both lagged prices and lagged output.

As others have noted, for a solution to the central banker's problem to exist, the coefficient on lagged output in the price equation, b , must have a real value. That is, a solution exists if and only if

$$(11) \quad \frac{1 - \lambda}{\lambda} \leq \frac{(1 - \beta\rho)^2}{4\rho^2\alpha^2\beta}.$$

Parkin (2000) points out that this condition is somewhat restrictive, since only large values of λ are consistent with high persistence in output (ρ close to one). This means that if λ is low and ρ is high, there is no solution. The reason is that under these circumstances the optimal response to stabilize output requires very high, even infinite, volatility of the price level (or inflation).¹⁰ Fortunately, most estimates that we know of suggest that central banks place much higher weight on inflation than they do on output volatility. For example, Cecchetti and Ehrmann (2002) estimate λ 's for a number of countries, and most of them are 3/4 or higher. So we view this problem as unlikely to occur in practice.

4.2.2 Society's Problem

With a complete characterization of the central bank's problem in hand, we can now turn to society's problem: what value of η should monetary policymakers be instructed to use? To figure this out, all we need to do is find the value of η that minimizes the social loss function, taking account of the central banker's behavior. Recall that we assume society minimizes a weighted average of inflation and output variability. We can write this as

$$(12) \quad L^S = \lambda \sigma_\pi^2 + (1 - \lambda) \sigma_y^2.$$

For now we look only at the case in which λ is the same for society and the central bank. Using the previous results, we can write this as

$$(13) \quad L^S = \left[D^2 \frac{2\lambda(1 - \rho)}{(1 + \eta)(1 - \rho\eta)} + (1 - \lambda) \right] \left[\frac{1}{1 - \rho^2} \left(\frac{c}{D} \right)^2 \right] \sigma_\varepsilon^2.$$

Taking the derivative with respect to η (noting that D is not a function of η and assuming that the condition [11] holds) yields the optimal hybrid-targeting regime:

$$(14) \quad \eta^* = \frac{1 - \rho}{2\rho}.$$

10. As we show in the appendix, this is a problem that only arises under discretion. If the central bank can be forced to commit to an instrument rule, then the problem always has a solution.

The result tells us that as ρ approaches 1, so that the shocks to output are extremely persistent, η^* goes to zero. As ρ shrinks, η^* grows, but we assume that it can never exceed 1. Importantly, the expression is consistent with Svensson's result. He shows that if one is restricted to choosing $\eta = 0$ or $\eta = 1$, then the threshold is at $\rho = 0.5$.

Before proceeding, we note that under commitment, where society can bind policymakers not just to an objective function but to an instrument rule as well, the best thing to do is to give the central bank society's loss function. That is not at all surprising. What is surprising is that if society's loss is in terms of the *price path* rather than inflation—that is, L^S is a function of σ_p^2 rather than σ_π^2 —then the discretionary solution is the same as the commitment solution.¹¹

4.2.3 Stabilization Bias

So far, we have been concerned with the benefits to be obtained from giving the central bank a hybrid target. But in addition to choosing η^* , society has the option of giving the central bank a λ that deviates from its own. The incentive for doing this comes from the fact that, left to their own devices, central bankers may choose to stabilize output more than is socially optimal. Avoiding this stabilization bias requires setting λ^{CB} above λ^S .

To see how this works, we return to equation (13) and note first that λ here represents social preferences and that D (defined in the previous section) is a function of the central bank's λ . Using this, we can rewrite the expression for the social loss as

$$(15) \quad L^S = \left[D(\lambda^{CB})^2 \frac{2(1-\rho)\lambda^S}{(1+\eta)(1-\rho\eta)} + (1-\lambda^S) \right] \cdot \left\{ \frac{1}{1-\rho^2} \left[\frac{1}{1+\alpha D(\lambda^{CB})} \right]^2 \right\} \sigma_\varepsilon^2.$$

This change has no impact on the degree of optimal hybrid targeting. η^* was not a function of λ before, and it is not now. But minimizing equation (15) requires not only finding η^* but also figuring out what λ^{CB} should be as well. The first-order condition for this second choice is given by

$$\frac{\alpha(1-\lambda^{CB})}{\lambda^{CB}} f(\rho) = \frac{\alpha(1-\lambda^S)}{\lambda^S},$$

where $f(\cdot)$ is an increasing function of ρ . So with given λ^S , as ρ rises, λ^{CB} rises as well.

Figure 4.1 plots the relationship between output persistence and λ^{CB}

11. If inflation's primary cost is that it makes long-term planning difficult, then this may be the case we should all be focusing on. See the appendix for details.

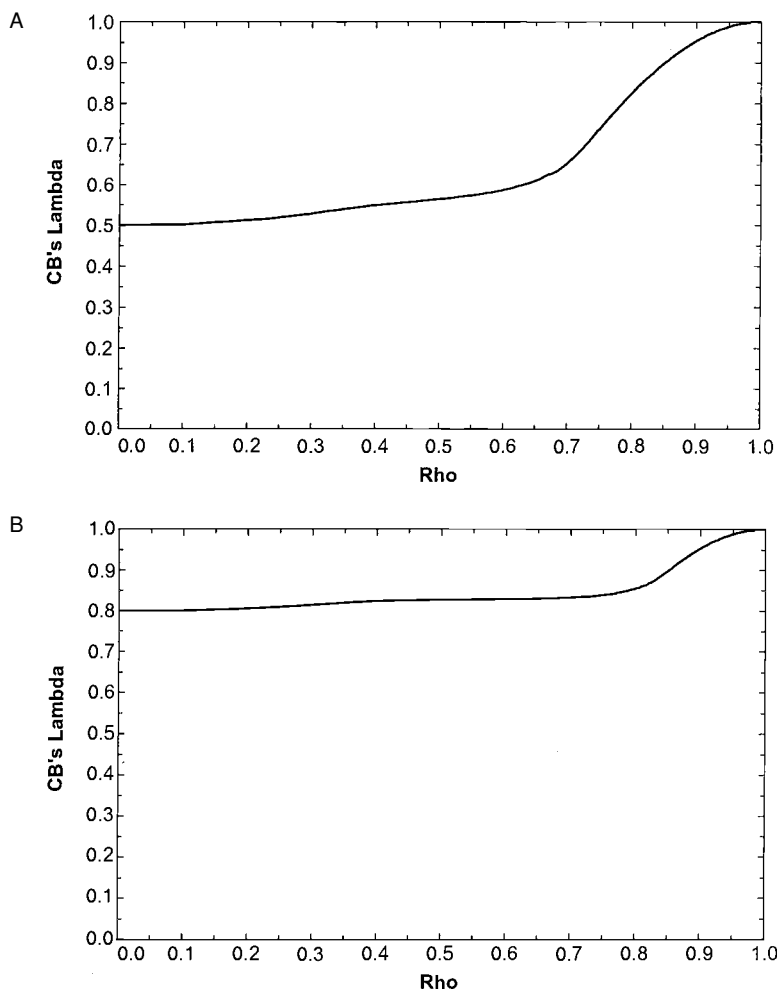


Fig. 4.1 Central bank's λ : *A*, Central bank's lambda when society's lambda = 0.5; *B*, Central bank's lambda when society's lambda = 0.8

when λ^s is 0.5 and 0.8. Throughout we assume that η is set at the optimal level, η^* in equation (14). The result is clear: the more persistent output is, the more conservative the central bank should be. And as the output approaches a random walk, the closer λ^{CB} gets to 1.

This is a good place to make another important point. In the last section we noted that there are times when the discretionary solution to the central banker's problem does not exist. Looking back at the restriction (11) required for existence, we see that there is always a solution when λ is big enough. So, if we are concerned that ρ may be high, we can avoid potential

difficulties by instructing the central banker to care almost exclusively about inflation.

4.3 Empirical Results

We now see that the optimal hybrid-targeting regime—the degree to which the central bank should target inflation relative to targeting the path of the price level—depends on how persistent output is. This leads us to ask the following questions: how persistent is output, and how close is the actual behavior of prices to what it would be under an optimal-targeting regime? The task of this section is to bring data to bear on these questions.

We do this in three steps. First, we estimate an empirical analog of the closed-economy model we studied in section 4.2. Second, since a number of countries we consider are small open economies, we introduce external factors into the estimation. Finally, we posit a social loss function in order to do welfare comparisons and measure the gains from adopting an optimal hybrid target.

4.3.1 Closed Economy

Our strategy is the following. Using quarterly data on consumer prices and industrial production, we estimate equations (9) and (10).¹² (The data are all described in the appendix.) Taking account of the serial correlation in output, we use the following specifications:

$$(16) \quad y_t = \rho y_{t-1} + \sum_{i=1}^4 \gamma_i \Delta y_{t-i} + e_{1t}$$

$$(17) \quad p_t = \eta p_{t-1} + b_1 y_{t-1} + b_2 y_{t-2} + b_3 y_{t-3} + b_4 y_{t-4} + e_{2t},$$

where y is computed as the deviation of log output from Hodrick-Prescott (HP) filtered output, and p measures the deviation of the log price level from a measure of the target. During the periods when countries were employing inflation targets, we used the target itself for this computation.¹³ In the absence of an inflation target, we used a Hodrick-Prescott filter.

The results for both the full sample (1980s and 1990s) and just the last decade are reported in tables 4.1 and 4.2. Estimates range widely.¹⁴ The first table shows estimates of ρ , together with standard errors. The important thing to notice is that ρ ranges from a low of 0.29 to a high of 0.82 and that

12. We note that our exact results are not invariant to the choice of the frequency of the data.

13. For the cases in which we have data for an explicit inflation target, we compute the price-path target as $p_t^* = \log(\text{CPI}_{t-1}) + \pi^*$, where π^* is the annual inflation target. Details are in the appendix.

14. All estimates throughout the paper are median-bias corrected using the empirical distributions that are also used to compute the standard errors.

Table 4.1 **Output persistence: The closed-economy case**

Country	Full sample		1990s	
	$\hat{\rho}$	Standard error	$\hat{\rho}$	Standard error
Australia	0.64	0.10	0.49	0.18
Austria	0.76	0.19	0.66	0.36
Canada	0.73	0.06	0.74	0.09
Chile	0.57	0.21	0.47	0.43
Denmark	0.56	0.14	0.31	0.23
Finland	0.78	0.07	0.65	0.13
France	0.61	0.15	0.61	0.15
Germany	0.70	0.10	0.61	0.17
Ireland	0.56	0.12	0.50	0.19
Israel	0.56	0.09	0.29	0.15
Italy	0.71	0.10	0.63	0.15
Japan	0.78	0.05	0.69	0.09
Korea	0.58	0.10	0.60	0.13
Mexico	0.64	0.10	0.69	0.15
The Netherlands	0.64	0.15	0.68	0.23
New Zealand	0.58	0.10	0.58	0.15
Norway	0.43	0.16	0.55	0.19
Portugal	0.76	0.08	0.69	0.14
Spain	0.72	0.07	0.70	0.11
Sweden	0.71	0.09	0.60	0.13
Switzerland	0.33	0.22	0.35	0.33
United Kingdom	0.80	0.07	0.78	0.08
United States	0.76	0.04	0.82	0.06

Notes: Estimates $\hat{\rho}$ are small sample bias-corrected autocorrelation coefficients from fourth-order autoregression using industrial production, equation (16). All data are quarterly data, seasonally adjusted and filtered using a Hodrick-Prescott filter. The full sample is 1980 Q1 to 2001 Q4 for non-euro-area countries. For countries in EMU, the sample ends in 1998 Q4. Standard errors are constructed from nonparametric bootstrap with 3,000 replications.

it is unstable across time periods. Both the range and instability have important implications for policy, so we will return to them later.¹⁵

Table 4.2 reports our estimates of the optimal hybrid-targeting regime, $\hat{\eta}^*$, as well as the estimate that is implied by the actual behavior of prices in each country, $\hat{\eta}$. Our estimates of ρ suggest that a number of countries should be putting significant weight on the price path, $\hat{\eta}^* \ll 1$, but virtually all of them exhibit behavior that is closer to inflation targeting, $\hat{\eta} \approx 1$. Given these estimates, we test whether $\hat{\eta} = \hat{\eta}^*$, and the answer is no. The p -value is reported in columns (3) and (6) of table 4.2.¹⁶

15. While we report results for a Hodrick-Prescott (HP) filter with parameter set to the standard 1600, experimentation in the range from 800 to 3200 leaves the character of our results unchanged.

16. Using a nonparametric bootstrap, we compute the empirical distribution of $\hat{\eta}^*$ and then report the p -value for $\hat{\eta}$ in that distribution.

Table 4.2 The optimal hybrid-targeting regime: The closed-economy case

Country	Full sample			1990s		
	$\hat{\eta}^*$	$\tilde{\eta}$	p -value testing $\tilde{\eta} = \eta^*$	$\hat{\eta}^*$	$\tilde{\eta}$	p -value testing $\tilde{\eta} = \eta^*$
Australia	0.29	0.81	0.01	0.50	0.69	0.33
Austria	0.15	0.68	0.07	0.25	0.46	0.26
Canada	0.18	0.94	0.00	0.18	0.88	0.00
Chile	0.36	0.72	0.22	0.31	0.64	0.31
Denmark	0.39	0.67	0.17	1.00	0.24	0.14
Finland	0.14	0.91	0.00	0.26	0.84	0.03
France	0.31	0.80	0.08	0.32	0.80	0.08
Germany	0.21	0.83	0.00	0.32	0.67	0.15
Ireland	0.39	0.81	0.09	0.49	0.63	0.37
Israel	0.39	0.90	0.03	1.00	0.79	0.26
Italy	0.21	0.94	0.00	0.29	1.00	0.05
Japan	0.14	0.75	0.00	0.23	0.70	0.01
Korea	0.35	0.90	0.02	0.34	0.56	0.19
Mexico	0.28	0.86	0.01	0.23	0.83	0.05
The Netherlands	0.28	0.88	0.06	0.24	0.60	0.17
New Zealand	0.36	0.93	0.03	0.36	0.54	0.25
Norway	0.64	0.77	0.39	0.38	0.56	0.32
Portugal	0.16	0.88	0.00	0.23	0.83	0.04
Spain	0.19	0.80	0.00	0.22	0.92	0.01
Sweden	0.21	0.84	0.00	0.33	0.53	0.21
Switzerland	1.00	0.89	0.44	0.89	0.87	0.36
United Kingdom	0.12	0.76	0.00	0.14	0.44	0.01
United States	0.16	1.00	0.00	0.11	0.91	0.00

Source: Data sources are all described in the appendix.

Notes: Estimates of $\hat{\eta}^*$ are constructed using the $\hat{\rho}$ in table 4.1. Estimates of $\tilde{\eta}$ are the coefficient on the lag of prices from equation (17). The p -values for the tests are constructed using a nonparametric bootstrap with 3,000 replications.

4.3.2 Open Economy

To take account of the fact that countries like Israel, Belgium, and Ireland are small and open, we introduce external factors into our analysis. Following Svensson (2000), we introduce import prices into the Phillips curve (6):

$$(18) \quad y_t = \rho y_{t-1} + \alpha(p_t - p_t^e) + \phi_y p_t^F + \varepsilon_t,$$

where p_t^F is the foreign price level denominated in domestic currency. With this modification, all of the results in section 4.2 go through, and we can rewrite empirical specification equations (9) and (10) as

$$(19) \quad y_t = \rho y_{t-1} + \sum_{i=1}^4 \gamma_i \Delta y_{t-i} + \phi_y p_t^F + e_{1t}$$

$$(20) \quad p_t = \eta p_{t-1} + b_1 y_{t-1} + b_2 y_{t-2} + b_3 y_{t-3} + b_4 y_{t-4} + \phi_p p_t^F + e_{2t}.$$

Table 4.3 **Output persistence: The open-economy case**

Country	Full sample		1990s	
	$\hat{\rho}$	Standard error	$\hat{\rho}$	Standard error
Australia	0.66	(0.09)	0.58	(0.19)
Austria	0.84	(0.23)	0.63	(0.50)
Canada	0.75	(0.06)	0.73	(0.11)
Chile	0.61	(0.07)	n.a.	n.a.
Denmark	0.61	(0.15)	0.13	(0.31)
Finland	0.79	(0.05)	0.78	(0.14)
France	n.a.	n.a.	0.61	(0.17)
Germany	0.73	(0.11)	0.69	(0.20)
Ireland	0.48	(0.14)	0.60	(0.22)
Israel	0.56	(0.09)	0.15	(0.17)
Italy	0.73	(0.09)	0.63	(0.13)
Japan	0.73	(0.04)	0.59	(0.09)
Korea	0.67	(0.11)	0.60	(0.12)
Mexico	0.67	(0.11)	0.53	(0.21)
The Netherlands	0.65	(0.17)	0.59	(0.29)
New Zealand	0.59	(0.10)	0.64	(0.16)
Norway	0.46	(0.16)	0.67	(0.21)
Portugal	0.78	(0.08)	0.75	(0.18)
Spain	0.73	(0.05)	0.74	(0.09)
Sweden	0.74	(0.10)	0.65	(0.16)
Switzerland	0.32	(0.24)	0.18	(0.43)
United Kingdom	0.81	(0.07)	0.79	(0.04)
United States	0.78	(0.03)	0.84	(0.04)

Source: See appendix and notes to table 4.1 for data sources.

Note: n.a. = not available.

Table 4.3 reports estimates of output persistence, ρ , after accounting for these external factors. The results are very similar to those in table 4.1. The correlation between these two sets of estimates is 0.96 for the full sample and 0.89 for the 1990s, and the mean absolute difference between the estimates is 0.03 and 0.075, respectively. Looking at the estimates of the various measures of η in table 4.4, our conclusions from the closed-economy analysis remain. In virtually every case, our estimate of the optimal hybrid target has η well below 1, closer to price-path targeting than inflation targeting, but the actual behavior of prices in these countries suggests something close to inflation targeting.

It is interesting to relate all of these results to what King (1999) referred to as an evaluation horizon for central bankers. He suggested that in practice an inflation-targeting central bank will be evaluated on whether it met its target on average over some number of years. The evaluation horizon is related to the hybrid regime. The longer the period over which inflation is averaged, the closer the regime is to price-path targeting. Using this intu-

Table 4.4 Optimal hybrid-targeting regime: The open-economy case

Country	Full sample			1990s		
	$\hat{\eta}^*$	$\tilde{\eta}$	p -value testing $\tilde{\eta} = \eta^*$	$\hat{\eta}^*$	$\tilde{\eta}$	p -value testing $\tilde{\eta} = \eta^*$
Australia	0.26	0.82	0.01	0.37	0.70	0.20
Austria	0.09	0.52	0.09	0.29	0.38	0.32
Canada	0.17	0.95	0.00	0.19	0.90	0.00
Chile	0.20	0.57	0.17	n.a.	n.a.	n.a.
Denmark	0.31	0.52	0.20	1.00	0.30	0.37
Finland	0.13	0.89	0.00	0.14	0.83	0.01
France	n.a.	n.a.	n.a.	0.33	0.67	0.17
Germany	0.19	0.72	0.01	0.21	0.75	0.09
Ireland	0.53	0.78	0.27	0.32	0.71	0.20
Israel	0.39	0.90	0.02	1.00	0.80	0.23
Italy	0.19	0.83	0.01	0.25	0.90	0.07
Japan	0.18	0.71	0.00	0.35	0.70	0.06
Korea	0.26	0.86	0.02	0.29	0.48	0.26
Mexico	0.26	0.88	0.01	0.42	0.88	0.20
The Netherlands	0.26	0.83	0.10	0.30	0.60	0.27
New Zealand	0.34	0.99	0.01	0.28	0.50	0.19
Norway	0.56	0.77	0.32	0.24	0.62	0.16
Portugal	0.14	0.85	0.00	0.16	0.85	0.04
Spain	0.18	0.78	0.00	0.17	0.90	0.01
Sweden	0.19	0.80	0.00	0.26	0.53	0.16
Switzerland	1.00	0.86	0.48	1.00	0.87	0.34
United Kingdom	0.11	0.76	0.00	0.13	0.44	0.00
United States	0.14	0.91	0.00	0.10	0.63	0.00

Source: See appendix and notes to text table 4.2.

Note: n.a. = not available.

ition, we can construct approximate measures of the horizon as $(1/\eta)$. For many countries we find that η^* is between 0.2 and 0.3, implying a horizon of between three and four quarters. To get a number that is usable in practice, we need to add another four to six quarters, the length of time that it takes for policy changes to have an impact on prices and output. The implication is that the evaluation horizon should be in the range of two to three years.

Before continuing, note that we recomputed all of the results for both the closed- and open-economy versions of our model substituting core consumer prices for the headline measures used in sections 4.3.1 and 4.3.2. Tables analogous to 4.2 and 4.4 are in the appendix. Overall, we find that the change in the price measure makes very little difference. Estimates of $\tilde{\eta}$ from the price equation are highly correlated between the two sets of matching results. For the full sample, the correlation for the seventeen countries for which we have data is 0.79 for the closed-economy model and 0.83 when import prices are included.

4.3.3 Loss Comparison

Simply computing the optimal value for η , the degree of a hybrid regime, is only the first step. What we really want to know is whether adopting the optimal hybrid makes any difference to welfare. Given the fact that estimates of η are fairly imprecise, this question is particularly important. To address it, we construct estimates of the social loss, L^S , for different targeting regimes and compare them. Computing the loss requires that we choose a series of parameters. Before turning to the data, it is useful to look at some simulations. Using the theoretical results, we can estimate the extent of the welfare gain that comes from going from an inflation-targeting regime to an optimal one. That is, we compare $L^S(\eta = 1)$ with $L^S(\eta = \eta^*)$ for various values of the parameters of the model. Note that throughout this exercise we assume that the preference parameter λ is the same for society and the central bank.

While it would be interesting to look across a wide range of values for the preference parameter λ , output persistence ρ , and the slope coefficient α , the condition (11) places restrictions on the relationship among these. So instead we look at a representative example. First, the restriction has a few simple properties: (a) given α , the higher ρ the higher the minimum λ ; and (b) given ρ , the higher α the higher the minimum λ . What that means is that the more persistent output and the flatter the aggregate supply curve—that is, the inverse of α in equation (6)—the higher the preference for inflation stability has to be for there to be a solution to the central bank's problem. To understand how restrictive this is, we have done a few simple calculations. Setting the discount factor $\beta = 0.99$, we see that for $\alpha = 0.5$ and $\rho = 0.7$, λ must be greater than 0.65. As α decreases, the range of permissible values grows. So when $\alpha = 0.3$, λ can be as low as 0.4 for $\rho = 0.7$. This creates a potential problem for the choice of α . While we would like to work with relatively low values, we choose $\alpha = 0.5$. This is the choice made by Dittmar, Gavin, and Kydland (1999), who use estimates in Rudebusch and Svensson (1999) as justification.

Using these parameter values, we examine the improvement in the social loss for each country for two changes: (a) moving from strict inflation targeting to the optimal hybrid regimes, that is, $L^S(\eta^*)/L^S(\eta = 1)$; and (b) shifting from a strict price-path targeting regime to the optimal hybrid, $L^S(\eta^*)/L^S(\eta = 0)$. Throughout we assume that the preference parameter $\lambda = 0.8$ and the discount rate $\beta = 0.99$. The results are somewhat sensitive to the choice of λ but not to the choice of β . Looking at table 4.5, we see that there is an important pattern. In no case does a move from price-path targeting to the optimal hybrid bring a sizable welfare gain. The same is not true of a move from inflation targeting. That is, the first and third columns include numbers that are far below 1—for example, 0.82 for Canada and 0.87 for Germany—while the second and fourth columns contain none.

Table 4.5 Loss comparison

Country	Full sample		1990s	
	$L^S(\eta^*)/L^S(\eta = 1)$	$L^S(\eta^*)/L^S(\eta = 0)$	$L^S(\eta^*)/L^S(\eta = 1)$	$L^S(\eta^*)/L^S(\eta = 0)$
Australia	0.94	0.99	0.97	0.99
Austria	0.71*	0.71*	0.95	1.00
Canada	0.82	1.00	0.86	1.00
Chile	0.91	0.99	0.71*	0.71*
Denmark	0.96	0.99	1.00	0.96
Finland	0.71*	0.71*	0.72	1.00
France	0.96	0.99	0.96	0.99
Germany	0.87	1.00	0.91	0.99
Ireland	0.99	0.99	0.96	0.99
Israel	0.98	0.99	1.00	0.96
Italy	0.87	1.00	0.94	0.99
Japan	0.85	1.00	0.97	0.99
Korea	0.93	0.99	0.95	0.99
Mexico	0.93	0.99	0.98	0.99
The Netherlands	0.96	0.99	0.97	0.99
New Zealand	0.96	0.99	0.95	0.99
Norway	0.99	0.99	0.93	0.99
Portugal	0.71	1.00	0.80	1.00
Spain	0.86	1.00	0.84	1.00
Sweden	0.86	1.00	0.94	0.99
Switzerland	1.00	0.97	1.00	0.98
United Kingdom	0.71*	0.71*	0.71*	0.71*
United States	0.71*	0.71*	0.71*	0.71*

Notes: Computations use $\alpha = 0.5$ and $\lambda = 0.8$, as well as the estimated value of ρ reported in table 4.1. Asterisks indicate values of (α, λ, ρ) for which restriction (11) is not met, and so the loss cannot be computed. The reported value is the minimum for which it can be computed.

It is worth examining this result in more detail. Figure 4.2 plots the two ratios $L^S(\eta^*)/L^S(\eta = 0)$ and $L^S(\eta^*)/L^S(\eta = 1)$ for a range of values for ρ and λ , assuming $\alpha = 0.5$ and $\beta = 0.99$. Taken together, these give us a striking picture of the potential benefits from adopting various regimes. First, note from panel A that even if ρ is very small, and so the optimal regime is close to one of pure inflation targeting, the loss from adopting price-path targeting is small. Only when λ is set to $2/3$, a relatively low value, and when output has virtually no persistence does a move from price-path targeting to the optimal hybrid imply a welfare gain of as much as 10 percent.

This is in stark contrast to panel B of figure 4.2, where we see the consequences of shifting from a pure inflation-targeting regime to the optimal hybrid. As output persistence rises above 0.6, the ratio of the losses starts to decrease very quickly. (Note that the lines end at the point where restriction [11] is no longer met.) That is, the gain from moving from inflation targeting to the optimal hybrid can be very large. To use Svensson's




Fig. 4.2 Loss comparing targeting regimes with optimal targeting: *A*, Comparing price-path targeting to the optimal hybrid regime ($\alpha = 0.5$); *B*, Comparing inflation targeting to the optimal hybrid regime ($\alpha = 0.5$)

terminology, there is a “free lunch,” and it can be big. And since we are unsure how big ρ really is, it is likely prudent to move to price-path targeting.

4.4 Conclusion

We have examined whether a country is well advised to target inflation, target the price path, or do something in between. The issue turns on the persistence of output deviations from their trend. With high persistence, which is what we tend to observe, our theoretical results suggest that countries are best off if they adopt a hybrid target that is close to price-path targeting. But such a policy regime would be difficult to adopt for two reasons. First, there is the technical one. The exact targeting procedure depends on the estimation of both the output trend and output persistence, both of which are going to be measured with substantial error. Second, the success or failure of any monetary policy regime rests critically on the ability of central bankers to communicate what they are doing to the public. Explaining a hybrid target would be challenging for even the best central bankers.

Taking these problems into account, we examine the welfare loss from adopting pure inflation or price-path targeting rather than the optimal hybrid. Our conclusion is that price-path targeting is less risky, in that the maximum social loss from being wrong—choosing price-path targeting when something else is better—is much smaller than if one chooses inflation targeting.

Appendix

Data Description

All data are quarterly beginning in quarter 1 (Q1) of 1980. For European Monetary Union (EMU) countries, data are through 1998 Q4. For non-EMU countries, data are through 2001 Q4.

1. Prices: Consumer Price Index (CPI) from the International Monetary Fund (IMF) International Financial Statistics.

2. Output: Industrial Production from the IMF International Financial Statistics (IFS), except for Portugal and Ireland, which are entirely from the Organization for Economic Cooperation and Development (OECD); New Zealand is from the OECD for 2000 Q3 on; Italy is from OECD for 2001 Q1 on; and Chile is manufacturing production only.

3. Core consumer prices: From the OECD.

4. Import prices: The import price index from the IMF International Financial Statistics, except for Spain, New Zealand, the Netherlands, Canada, France, Ireland, Israel, Italy (where the unit value of imports

from IFS is used), Mexico (Import Price Index from Haver Analytics), Austria (which uses the German CPI), and Portugal (an equally weighted average of the CPIs for the United Kingdom, Spain, France, and Germany).

5. Inflation targets are computed from the Inflation Targeting Country Fact Sheets” by Frank Gaenssmantel of the Institute of International Economics, courtesy of Edwin Truman.

The inflation targets are listed in table 4A.1. The target p_t^* is computed as follows:

$$(A1) \quad p_t^* = p_{t-1} + \pi^*$$

$$(A2) \quad p_t = \log(\text{CPI}_t) - p_t^* \quad \text{when there is } p_t^*$$

$$(A3) \quad p_t = \log(\text{CPI}_t) - \text{HPtrend}_t \quad \text{otherwise,}$$

where π^* is the annual inflation target in table 4A.1, divided by four. When the target is a range, the midpoint is used.

The Commitment Case

Our solutions in the text assume that the central bank operates under discretion. Discretion means that policymakers reoptimize the loss function every period after observing the state variable y_{t-1} and the shock ε_t . The alternative to this is commitment, in which the central bank optimizes once and commits to an instrument rule once and for all.

To find the commitment solution we take the derivative of the central bank’s loss in equation (5) with respect to p_t and p_t^c , subject to the constraint imposed by the Phillips curve in equation (6). The resulting policy rule, the equivalent to equation (8), is

$$(A4) \quad p_t = \eta p_{t-1} + \tilde{c} \varepsilon_t,$$

where

$$\tilde{c} = -\frac{\tilde{D}}{1 + \alpha \tilde{D}} \quad \text{and} \quad \tilde{D} = \frac{\alpha(1 - \lambda)}{\lambda(1 - \rho^2 \beta)}.$$

This is exactly the same as the case under discretion considered in section 4.2.1, except that $b = 0$. That is, under commitment the optimal response is to react only to the past price level and the shock, not to y_{t-1} . Recall, moreover, that the condition for a solution to exist under discretion, shown in equation (11), arises in computing b , and so it is not present here.

Continuing with the problem under commitment, society’s loss, the equivalent to equation (13), is now

$$(A5) \quad \tilde{L}^S = \left[\frac{2\lambda \tilde{c}^2}{1 + \eta} + \frac{1 - \lambda}{1 - \rho^2} \left(\frac{\tilde{c}}{\tilde{D}} \right)^2 \right] \sigma_\varepsilon^2.$$

Table 4A.1 Annual inflation target

Country	Period	Inflation target (%)
Australia	1993 Q1–2001 Q4	2.5
Austria	1993 Q1–2001 Q4	2.0
Canada	1992 Q1–1994 Q4	4.0
	1995 Q1–2001 Q4	3.0
Chile	1991 Q1–1991 Q4	18.0
	1992 Q1–1992 Q4	17.5
	1993 Q1–1993 Q4	11.0
	1994 Q1–1994 Q4	10.0
	1995 Q1–1995 Q4	8.0
	1996 Q1–1996 Q4	7.0
	1997 Q1–1997 Q4	6.0
	1998 Q1–1998 Q4	5.0
	1999 Q1–1999 Q4	4.3
	2000 Q1–2000 Q4	3.5
	2001 Q1–2001 Q4	3.0
Finland	1993 Q1–2001 Q4	2.0
Israel	1992 Q1–1992 Q4	14.5
	1993 Q1–1993 Q4	10.0
	1994 Q1–1994 Q4	8.0
	1995 Q1–1995 Q4	9.5
	1996 Q1–1996 Q4	9.0
	1997 Q1–1997 Q4	8.5
	1998 Q1–1999 Q4	4.0
	2000 Q1–2000 Q4	3.5
	2001 Q1–2001 Q4	3.0
Korea	1999 Q1–1999 Q4	3.75
	2000 Q1–2000 Q4	2.5
	2001 Q1–2001 Q4	3.0
Mexico	1995 Q1–1995 Q4	19.0
	1996 Q1–1996 Q4	20.5
	1997 Q1–1997 Q4	15.0
	1998 Q1–1998 Q4	12.0
	1999 Q1–1999 Q4	13.0
	2000 Q1–2000 Q4	10.0
	2001 Q1–2001 Q4	6.5
New Zealand	1992 Q1–1996 Q4	1.0
	1997 Q1–2001 Q4	1.5
Norway	2001 Q1–2001 Q4	2.5
United Kingdom	1992 Q1–2001 Q4	2.5

Source: Gaenssmantel (2002).

The η that minimizes this loss is trivially 1, which implies inflation targeting. Under commitment, it is optimal to simply give the central bank society's loss function.

When Society Prefers Price-Path Targeting

What if society's preferences are in terms of the path of the price level rather than an inflation target? In this case, the central bank's problem is

the same as the one in section 4.2.1. It is the social loss, equation (12), that changes. Assuming society cares about the price path implies that the social loss function is

$$(A6) \quad L^S = \lambda\sigma_p^2 + (1 - \lambda)\sigma_y^2.$$

Substituting in the solution for the central bank's problem, this becomes

$$(A7) \quad L^S = \left[\lambda D^2 \left(\frac{1}{1 - \eta^2} \right) \left(\frac{1 + \eta\rho}{1 - \eta\rho} \right) + (1 - \lambda) \right] \sigma_y^2.$$

Equation (A7) is the equivalent to text equation (13). The optimal η that minimizes this loss is zero. So, if society cares about the price path, then the central bank should be told to care about it, too.

Substituting Core for Headline Consumer Prices

The following tables are from substituting measures of the core CPI for the headline CPI in the computations of section 4.3. Table 4A.2 is the analog to text table 4.2, and table 4A.3 is the analog to text table 4.4. Note that

Table 4A.2 **Optimal hybrid-targeting regime: Closed economy with core CPI**

Country	Full sample			1990s		
	$\hat{\eta}^*$	$\tilde{\eta}$	<i>p</i> -value testing	$\hat{\eta}^*$	$\tilde{\eta}$	<i>p</i> -value testing
			$\tilde{\eta} = \eta^*$			$\tilde{\eta} = \eta^*$
Australia	n.a.	n.a.	n.a.	0.50	0.70	0.32
Austria	0.15	0.44	0.14	0.25	0.27	0.41
Canada	0.18	0.93	0.00	0.18	0.87	0.00
Chile	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.
Denmark	0.39	0.56	0.25	1.00	0.21	0.12
Finland	0.14	0.91	0.00	0.26	0.73	0.05
France	0.31	0.92	0.06	0.32	0.92	0.06
Germany	0.21	0.90	0.00	0.32	0.81	0.11
Ireland	0.39	0.81	0.08	0.49	0.67	0.33
Israel	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.
Italy	0.21	0.94	0.00	0.29	1.03	0.04
Japan	0.14	0.39	0.00	0.23	0.47	0.04
Korea	n.a.	n.a.	n.a.	0.34	0.63	0.16
Mexico	0.28	0.91	0.01	0.23	0.84	0.05
The Netherlands	0.28	0.79	0.07	0.24	0.52	0.20
New Zealand	0.36	0.81	0.03	0.36	0.54	0.26
Norway	n.a.	n.a.	n.a.	0.38	0.63	0.28
Portugal	n.a.	n.a.	n.a.	0.23	0.79	0.04
Spain	0.19	0.79	0.00	0.22	0.83	0.01
Sweden	0.21	0.78	0.00	0.33	0.46	0.27
Switzerland	1.00	1.02	0.39	1.00	0.93	0.34
United Kingdom	0.12	0.70	0.00	0.14	0.35	0.02
United States	0.16	0.97	0.00	0.11	0.68	0.00

Source: See appendix and notes to text table 4.2.
Note: n.a. = not available.

Table 4A.3 Optimal hybrid-targeting regime: Open economy with core CPI

Country	Full sample			1990s		
	$\hat{\eta}^*$	$\tilde{\eta}$	<i>p</i> -value testing	$\hat{\eta}^*$	$\tilde{\eta}$	<i>p</i> -value testing
			$\tilde{\eta} = \eta^*$			$\tilde{\eta} = \eta^*$
Australia	n.a.	n.a.	n.a.	0.37	0.72	0.21
Austria	0.09	0.26	0.23	0.29	0.17	0.45
Canada	0.17	0.94	0.00	0.19	0.88	0.00
Chile	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.
Denmark	0.31	0.44	0.29	1.00	0.05	0.35
Finland	0.13	0.91	0.00	0.14	0.70	0.01
France	n.a.	n.a.	n.a.	0.33	0.93	0.08
Germany	0.19	0.87	0.00	0.21	0.86	0.07
Ireland	0.53	0.79	0.24	0.32	0.58	0.25
Israel	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.
Italy	0.19	0.83	0.01	0.25	0.93	0.02
Japan	0.18	0.37	0.01	0.35	0.45	0.26
Korea	n.a.	n.a.	n.a.	0.29	0.62	0.15
Mexico	0.26	0.95	0.01	0.42	0.89	0.20
The Netherlands	0.26	0.79	0.09	0.30	0.42	0.37
New Zealand	0.34	0.82	0.02	0.28	0.51	0.19
Norway	n.a.	n.a.	n.a.	0.24	0.69	0.14
Portugal	n.a.	n.a.	n.a.	0.16	0.75	0.05
Spain	0.18	0.79	0.00	0.17	0.77	0.00
Sweden	0.19	0.76	0.00	0.26	0.46	0.21
Switzerland	1.00	1.02	0.41	1.00	0.94	0.33
United Kingdom	0.11	0.71	0.00	0.13	0.34	0.01
United States	0.14	0.94	0.00	0.10	0.69	0.00

Source: See appendix and notes to text table 4.4.

Note: n.a. = not available.

since the output equations (16) and (19) do not include the price level, the estimates of ρ and η^* are unchanged, and so the corresponding columns in the tables are identical. Comparing these results to those in the text, we conclude that substituting core for headline prices changes little.

References

- Batini, Nicoletta, and Anthony Yates. 2003. Hybrid inflation and price level targeting. *Journal of Money, Credit and Banking* 35 (3): 283–300.
- Cecchetti, Stephen G., and Michael Ehrmann. 2002. Does inflation targeting increase output volatility? An international comparison of policymakers' preferences and outcomes. In *Monetary Policy: Rules and Transmission Mechanisms*, Proceedings of the fourth annual conference of the Central Bank of Chile, ed. Norman Loayza and Klaus Schmidt-Hebbel, 247–74. Santiago, Chile: Central Bank of Chile.

- Cecchetti, Stephen G., Alfonso Flores-Lagunes, and Stefan Krause. 2002. Has monetary policy become more efficient? Brandeis University, International Business School, University of Arizona, Department of Economics, and Emory University, Department of Economics. Unpublished manuscript.
- Clark, Peter B., Charles A. E. Goodhart, and Haizhou Huang. 1999. Optimal monetary policy rules in a rational expectations model of the Phillips curve. *Journal of Monetary Economics* 43:497–520.
- Dittmar, Robert, and William T. Gavin. 2000. What do New Keynesian Phillips curves imply for price-level targeting? *Economic Review of the Federal Reserve Bank of St. Louis* 82 (April): 21–30.
- Dittmar, Robert, William T. Gavin, and Finn Kydland. 1999. The inflation-output variability trade-off and price-level targeting. *Economic Review of the Federal Reserve Bank of St. Louis* 81 (1): 23–39.
- Gaenssmantel, Frank. 2002. Inflation targeting country fact sheets. Institute of International Economics. Mimeograph.
- King, Mervyn. 1999. Challenges for monetary policy: New and old. In *New challenges for monetary policy*, proceedings of the Federal Reserve Bank of Kansas City symposium, 11–57. Kansas City, Mo.: Federal Reserve Bank of Kansas City.
- Mankiw, N. Gregory, and Ricardo Reis. 2001. Sticky information vs. sticky prices: A proposal to replace the New Keynesian Phillips curve. NBER Working Paper no. 8290. Cambridge, Mass.: National Bureau of Economic Research, May.
- Mishkin, Frederic S. Forthcoming. Inflation targeting. In *Encyclopedia of macroeconomics*, ed. Howard Vane and Brian Snowdon. London: Edward Elgar.
- Parkin, Michael. 2000. A note on inflation targeting versus price level targeting. University of Western Ontario, Department of Economics. Mimeograph.
- Rogoff, Kenneth. 1985. The optimal degree of commitment to an intermediate monetary target. *Quarterly Journal of Economics* 100:1169–90.
- Rudebusch, Glenn D., and Lars E. O. Svensson. 1999. Policy rules for inflation targeting. In *Monetary policy rules*, ed. John B. Taylor, 203–46. Chicago: University of Chicago Press.
- Söderlind, Paul. 1999. Solution and estimation of RE macromodels with optimal policy. *European Economic Review* 43:813–23.
- Svensson, Lars E. O. 1999. Price level targeting vs. inflation targeting. *Journal of Money, Credit and Banking* 31 (August): 277–95.
- . 2000. Open-economy inflation targeting. *Journal of International Economics* 50 (February): 155–83.
- Vestin, David. 2000. Price level targeting vs. inflation targeting in a forward-looking model. Working Paper no. 106. Stockholm, Sweden: Sveriges Riksbank, May.

Comment N. Gregory Mankiw

I like the starting point of this paper—the question of whether inflation targeting or price-level targeting is the better policy for a central bank to adopt. Like the authors, I think this is an important and still open question

N. Gregory Mankiw is currently chairman of the Council of Economic Advisers, Allie S. Freed Professor of Economics at Harvard University, and a research associate of the National Bureau of Economic Research.

in the analysis of monetary policy. And I agree with the paper's conclusion that, given our current understanding of the issue, price-level targeting is probably the better of the two alternatives. (For my approach to this issue, see Ball, Mankiw, and Reis 2003.)

The job of discussant, however, is like that of Mark Antony—not to praise the authors but to bury them. So even though there is a lot in this paper that I agree with, in my comments I will emphasize the points of disagreement. Going from the starting point (which I like) to the conclusion (with which I concur), this paper takes a few wrong turns along the way. Sometimes these wrong turns follow in the footsteps of the literature; other times the authors strike out in a mistaken direction all on their own.

Fortunately, one wrong turn the authors avoided is the use of an implausible model of inflation-output dynamics, although they make the mistake of apologizing for this fact. In a footnote, they say that they use a neo-classical Phillips curve for its tractability, suggesting that they would have preferred to use a New Keynesian Phillips curve. In my view, this gets things exactly backward. I think that every paper in this conference that uses the New Keynesian Phillips curve should apologize. Let me suggest the following footnote for those papers: "We use the New Keynesian Phillips curve even though its predictions about monetary policy are inconsistent with what most empirical studies find and with what every central banker knows to be true. We use this model because we think it is neat, and because that's what everybody else is doing."

Another way in which the authors' views differ from mine is in their acceptance of the Svenssonian approach to the analysis of monetary policy. As I understand it, the Svenssonian approach is based on the idea that two wrongs make a right, as least if the wrongs are well chosen. That is, society has a problem because monetary policy is made by discretion and thus suffers from time inconsistency. We can fix this problem, however, by assigning the central bank an objective function that differs from the true social welfare function. The Svensson insight is that the wrong of having an incorrect objective function can offset the wrong of having discretionary policy. This is a classic second-best type of analysis.

What puzzles me about this approach is the question of implementation. That is, how are we supposed to give the central bank this new objective function?

One possibility is that the central bank takes direction from a higher authority, such as Congress. In this case, why would the higher authority direct the central bank to have the wrong objective function? It seems more natural to direct the central bank to follow the optimal rule based on the true social welfare function. This is roughly McCallum's "just do it" viewpoint, and it is similar to the approach envisioned by Woodford's "timeless" perspective on monetary policy analysis. It is not at all obvious to me

why one type of direction from a higher authority to a central bank is more feasible than the other. That is, if the higher authority can assign an objective function to the central bank, it should be able to assign the central bank the constraint of being time consistent.

There is, however, another way to think about implementing the Svenssonian objective function. When we appoint central bankers, we can look around the population of candidates and pick someone who happened to think that the true social welfare function was the one that Svensson derives as the right one for a central bank to maximize. This central banker would be wrong, but he would be wrong in a useful way. This is akin to Rogoff's analysis of why we might want central bankers to be more conservative, in the sense of more inflation averse, than the general public.

This approach to implementation also strikes me as a bit odd. If a potential central banker is misguided about the social welfare function, why would that be the case? Most likely, he has the wrong model of the economy. The Svensson-Rogoff assumption is that the central banker has the right model but the wrong social welfare function. My experience is that people who are confused about one thing are often confused about other things as well. Looking for public servants who are confused in just the right way to offset the problems of discretionary policy does not seem like a winning strategy.

These comments, however, are aimed more at the broader literature than at this particular paper. Let me now put these larger concerns aside and turn to some issues that are more specific to this paper.

The empirical heart of this paper concerns the persistence of output. In Svensson's model, the desirability of inflation or price-level targeting depends on the autoregressive parameter in the output equation. The more persistent output shocks are, the more attractive price-level targeting becomes. An autoregressive parameter of 0.5 is a crucial cutoff.

This raises a natural question: how long is a period in the model? That is, what frequency of data should we use to implement the model? If output is AR(1) with parameter 0.8 in quarterly data, it is AR(1) with parameter 0.41 ($=0.8^4$) in annual data. We would reach a different conclusion about policy if we applied the model at an annual rather than a quarterly frequency. The Cecchetti-Kim paper uses quarterly data, but it does not explain why this is the right choice.

There is, however, something in the model that can be used to pin down the choice of data frequency: the timing of expectations. The Phillips curve in this model is based on one-period-ahead expectations of the price level. The paper does not tell us precisely how this equation is motivated, but one common approach is Fischerian labor contracts predetermining the nominal wage. In this case, the relevant issue is how far in advance wages are set. If wages are set one year ahead, rather than one quarter, applying the

model at an annual frequency would make more sense. In this case, the estimated autoregressive parameters would be much lower than those presented in the paper. By using higher-frequency data, the authors build in a bias toward their conclusion of price-level targeting.

Another nuts-and-bolts empirical issue that is crucial for this paper is the classic topic of detrending. Cecchetti and Kim look at persistence in quarterly output detrended with the Hodrick-Prescott (HP) filter. But an arbitrary parameter in this filter governs how much of the low-frequency movement in the data is filtered out. Their results in table 4.1 tell us that shocks to U.S. GDP have a half-life of about three quarters. I suspect that this result is more an artifact of the filter than a fact about the data. If they altered the smoothing parameter in the HP filter, they would be likely to get very different estimates for this key autoregressive parameter.

A central issue in this model is the length of time with which monetary shocks influence output. If we knew the answer to this question, we could use it to calibrate the autoregressive parameter and judge whether the parameter is bigger or smaller than $1/2$, the key cutoff. If I am right that we should be thinking at an annual frequency, because labor contracts are annual more often than quarterly, then the key question is this: does a monetary shock's effect on output dissipate by more or less than 50 percent if measured a year after the shock? I don't know the answer, but I doubt we can learn it from running univariate autoregressions using HP filtered output.

In closing, let me briefly address the big question of whether price-level targeting really would be a good monetary policy. There are now a lot of academic studies suggesting that it would be a good policy for a variety of reasons. My experience from talking to central bankers, however, is that they are often horrified at the idea. They have trouble imagining that a period of higher-than-target inflation should be followed by a period of lower-than-target inflation.

The reason they are horrified by this prospect, I think, is that in their hearts they don't really believe the Lucas critique. They tend to view the world through the lens of an expectations-augmented Phillips curve with adaptive expectations. If that model were truly structural, then price-level targeting would not be very attractive. Academics, however, are more likely to view that reduced form as an artifact of the monetary regime we have had over the past several decades. The reduced-form Phillips curve would look very different if a central bank adopted price-level targeting.

We academics, however, should be careful to maintain a bit of humility when we engage in this policy debate. We have to admit that our understanding of inflation-output dynamics is still primitive. Until we reach a consensus about the right model about the Phillips curve, we cannot be confident about the effect of any alternative monetary policy, especially proposals as radical as price-level targeting.

Reference

Ball, Laurence, N. Gregory Mankiw, and Ricardo Reis. 2003. Monetary policy for inattentive economies. NBER Working Paper no. 9491. Cambridge, Mass.: National Bureau of Economic Research.

Discussion Summary

In response to Gregory Mankiw's comments, *Lars Svensson* defended the relevance and importance of delegating an objective other than social welfare to the central bank. Social welfare is obviously too complex and multidimensional an objective to be operational for monetary policy, and assigning it to monetary policy is counterproductive, as monetary history has clearly shown. A large part of successful monetary policy reforms in many countries has instead consisted of assigning, by legislation or government instructions, simple and verifiable objectives for central banks, such as price stability or flexible inflation targeting, with (for instance) the understanding that the explicit or implicit output target is the natural output level rather than the socially optimal output level. This has resulted in better outcomes from a social welfare point of view. Other economic policy than monetary policy is then assigned to raise the natural output level toward the optimal level.

George Evans suggested that the New Keynesian Phillips curve was attractive because it is forward-looking, an assumption that policymakers would find plausible.

Bennett McCallum argued that the assumption that central bankers act in a discretionary manner was at odds with the assumption that they would be willing to minimize a delegated loss function that was different from their own. Concerning price-level targeting, in his earlier work on targeting the level of nominal GDP versus targeting nominal GDP growth, he had found beneficial effects of using growth but giving some small weight also to the level and hence inducing trend stationarity of nominal GDP.

Michael Woodford agreed with Gregory Mankiw that the assignment of an objective to the central bank was a problematic way of avoiding the losses caused by discretionary optimization. He argued, however, that price-level targeting was valuable independent of whether it was implemented by delegation of a loss function or in some other way—for example, because it was easily interpretable and robust across different model specifications.

Mervyn King expressed the view that price-level targeting would be too costly if the Phillips curve was backward looking, and that therefore the private sector's expectations formation was a key issue to judge whether

price-level targeting was desirable. He also argued that price-level targeting may give the impression that the central bank was too much engaged in fine-tuning the economy, unless the horizon chosen over which to return the price level to its target was sufficiently long.

John Berry expressed the concern that the choice of the target path for the price level had to take into consideration measurement error in the price index in terms of which the target was formulated.

Stephen Cecchetti replied that he agreed that the delegation of an objective function was an idea that in practice would be difficult to implement.