When Is Price-Level Targeting a Good Idea?

Richard Barnett and Merwan Engineer

Introduction

This paper critically examines a new literature that supports price-level targeting (PT) over inflation targeting (IT).[1] This literature shows that it is better for the central bank (the Bank) to gradually wrestle the price level, rather than the inflation rate, back to the target path. PT provides a firm nominal anchor for expectations, and thereby conditions private sector expectations in a way that reduces inflation variability and improves welfare. These recent results appear to contradict the conventional view that PT is a "bad idea" because it increases the variability of inflation and output relative to IT. In this paper, we attempt to sort out these differing views within a standard theoretical framework in which the Bank minimizes a quadratic loss function subject to a log-linear Phillips curve (PC).

1. By PT we mean any policy that stabilizes the price-level index around a price-level-target path in the long run. The policy may or may not explicitly use a "PT rule," but it precludes long-run, price-level drift from the implicit target path. As PT stabilizes inflation, we distinguish IT as any policy that stabilizes inflation but implies price-level drift. See section 1.1.

The only country to adopt PT policy was Sweden in the 1930s (see Berg and Jonung 1999). Bernanke et al. (1999) and Mishkin (2001) evaluate the international experience with IT policy. Bernanke et al. describe the practice of IT as a framework for policy analysis within which "constrained discretion" can be exercised. The framework can serve two important functions: (i) communication between the policy-maker and the public, and (ii) increased discipline and accountability for monetary policy.

The conventional case for PT is that it facilitates long-term planning and nominal contracting.[2] Because it precludes price-level drift, PT avoids the arbitrary transfer of wealth and the associated increased debt burdens from deflation that might precipitate widespread bankruptcy. A policy of targeting a fixed price level is particularly appealing. Money becomes a unit of account and standard for deferred payment with constant value. Nominal values become real values, and this reduces calculation and menu costs and improves the role of prices in allocating resources. Targeting a fixed price level also eliminates the distortions from the nominal tax system.

Nevertheless, "conventional wisdom" has been skeptical of PT. On the benefits side, it can be argued that fiscal instruments are the appropriate means for solving taxation and distribution problems. Fischer (1994) argues that the benefits of more stable long-term nominal contracting are not likely to be substantial given that other means (e.g., indexed bonds, contingent contracts) exist to ameliorate long-run price uncertainty. From a practical standpoint, McCallum (1999) argues that the net reduction in price uncertainty in the United States under PT would be small.

The main argument against PT on the costs side is that it induces both higher short-run inflation and output variability than does IT.[3] Shocks that move the price level above (below) the target path lead the monetary authority to disinflate (inflate) with lower (higher) than average inflation to move towards the target path. This whipsawing of inflation above and then below trend induces short-run inflation variability relative to IT, since IT allows price-level drift and aims for only the target inflation rate. Greater inflation variability induces greater output variability along the short-run PC. Furthermore, when nominal rigidities are entrenched, as they might be at around zero inflation, the fear is that a policy that involves deflating might induce a recession. Fischer (1994, 282) writes, "Price level targeting is thus a bad idea, one that would add unnecessary short-term fluctuations to the economy."

This conventional argument has been challenged in a series of recent papers that show that PT is welfare-improving. There are two strands to this literature, corresponding to whether or not the Bank can commit to future policy. The strongest case for PT is when the Bank can commit to policy, and expectations are forward-looking in a New-Keynesian PC. Then a price-level-trend stationary policy is optimal, using a standard social loss function

---

2. Black, Coletti, and Monnier (1998), Duguay (1994), Feldstein (1997), and Konieczny (1994) survey the benefits of price-level stability.

3. Svensson (1999d) calls this the "conventional wisdom" and cites Fischer (1994), Haldane and Salmon (1995), and Lebow, Roberts, and Stockton (1992). Fillion and Tetlow (1994) find that PT increases output variability but decreases inflationary variability.

defined on inflation and output.[4] The price-level target forms a firm nominal anchor for expectations. Agents know that shocks that move the price level above trend will eventually be countered by measures by the Bank to move the price level back to the target path. Because they know this involves inflation below trend in the medium term, agents lower their expectations of inflation, thus reducing inflation and inflation variability and increasing welfare.[5]

The second strand of the literature considers a Bank that cannot commit to future policy. Assigned the social loss function, the Bank is unable to condition private sector expectations in a desirable way. This is because the social loss function has an inflation-targeting objective (ITO), so that the Bank's dynamic program ignores the history prior to the current period. The Bank does not persist in battling past inflation; rather, it looks forward and engages in IT.

In an intriguing argument, Svensson (1999d) shows that assigning a Bank a loss function with a price-level-targeting objective (PTO) yields a PT policy that may reduce inflation variability without affecting output variability. This "free-lunch" result depends on substantial endogenous output persistence in the New-Classical PC.[6] Dittmar and Gavin (2000) and Vestin (2000) extend this analysis to the case where expectations are forward-looking in a New-Keynesian PC. They show that the free-lunch argument applies without the need for persistence terms in the PC. Thus, assigning the Bank a PTO appears to improve welfare if expectations are forward-looking or if there is substantial endogenous persistence.

We will examine the robustness of the new arguments for PT in the framework in which they were developed, where the Bank minimizes a quadratic loss function subject to a log-linear PC. This "linear-regulator" approach has recently been put on a firm microeconomic footing. It is also the standard framework used for much of applied monetary policy analysis. For these reasons, we consider it appropriate for this first basic primer on PT.

---

4. Clarida et al. (1999), Woodford (1999a, 1999c), and Vestin (2000) demonstrate this result. Also see Backus and Driffill (1986) and Currie and Levine (1993).

5. Black, Macklem, and Rose (1998) find that the inflation and output variability often decrease when the price level is weighted in the policy rule. Unlike earlier studies (see footnote 3), they allow expectations to adjust to take account of price-level control.

6. Svensson (1999d) terms this a "free lunch" because only inflation variability falls. However, the result is stronger because it is often possible to improve both output and inflation variability by assigning the Bank a loss function with a different weight than Svensson uses.

Within the linear-regulator framework, the results for and against PT depend on whether: (i) the Bank can commit to policy or acts under discretion, (ii) expectations in the PC are forward-looking or predetermined, (iii) output persistence in the PC is endogenous or exogenous, (iv) the Bank can target current or forecast variables, and (v) the Bank can be assigned alternative loss functions. In the body of the paper, we detail the role of these factors. Our answer to the title's question—when is price-level targeting a good idea?—is found in the concluding section.

The paper proceeds as follows. Section 1 details the use of language and provides examples of specific target rules. Section 2 outlines the ingredients in the basic models. Section 3 describes the Bank's problem when it can credibly commit to future policy. Results are presented for a hybrid PC when current variables can be controlled, as well as when there is a lag in controlling target variables. The Bank's problem under discretion is described in section 4. Section 5 describes how the results would change under different social welfare functions. Section 6 answers the title's question and briefly discusses issues not covered in our main analysis.

1 Terms and Examples

1.1 Targets and targeting

PT presupposes a price-level target as an important long-run goal for monetary policy. Most often, the target is thought of as a fixed price level, $p^* = p_t^* = p_{t-1}^*$, where prices are expressed in logarithms. More generally, it can be considered a target path for the price level,

$$\{p_t^*\}_{t=t_0}^{\infty},$$

where t_0 is the beginning period. Similarly, IT presupposes an inflation target as an important long-run goal. We take the price-level- or inflation-target path as exogenous and focus on the consequences of targeting policy.

The terms "inflation targeting" and "price-level targeting" are used loosely in the literature to refer to the long-run implications of a policy, as well as to its immediate intention. This is confusing because, as we shall see, a stated aim of targeting inflation can have the effect of PT. To avoid this confusion, we define PT and IT so that they are mutually exclusive. Targeting relates to the long-run implications of a policy. Targeting rules are later defined relative to the stated intentions of a policy.

Definition: *Price-level targeting* (PT) is a policy that has the effect of systematically responding to deviations of the price level from the price-level-target path to preclude long-run price-level drift.

Definition: *Inflation targeting* (IT) is a policy that systematically responds to deviations of the inflation rate from the target inflation rate in a way that implies long-run price-level drift.

PT is any policy that successfully acts to stabilize the price level around the target path. With a fixed price-level target, PT is price-stationary and also inflation-stationary around zero inflation. When the target path grows at a constant rate, $\pi^* > 0$, PT policy is price-level-trend stationary and hence inflation-stationary around π^*. We show below that PT also implies that average inflation is stationary around π^*. Thus, PT and "average inflation targeting" are closely related, because both the price level and the average inflation rate incorporate past inflation rates.

IT is a policy that successfully acts to stabilize the inflation rate around the inflation target path. For a fixed inflation target, π^*, IT is inflation-stationary around π^*. But to distinguish IT from PT, we restrict use of the term IT to policies that involve price-level drift so that they are not also price-stationary. In this way, there is no overlap and confusion in the use of the terms.

The policies achieve different ends. Consider PT with a fixed price-level target, $p^* = 0$ (in logs), versus IT with a zero-inflation target, $\pi^* = 0$. PT yields price-level stability around p^* and zero inflation. IT also yields zero inflation but does not deliver price-level stability around p^* because of price-level drift.[7] This case is illustrated in Figure 1 (section 3). Series 1 and 2 describe IT policies, whereas series 3 is a PT policy. In response to a 1 per cent inflation shock in period 1, series 3 produces a deflation response in subsequent periods until the original price level, $p^* = 0$, is restored. Although series 2 corrects towards $p^* = 0$, it never gets there and, hence, displays price-level drift.

1.2 Target rules

We adopt the following terms from Svensson and Woodford (1999).

Definition: A *specific targeting rule* gives a formula for how target variables relate to targets.

Definition: A *general targeting rule* uses a Bank objective function, constraints, and an optimization routine to describe how target variables relate to the target levels.

7. When the inflation target can be hit in every period, there is no price-level drift. We refer to this as PT. A meaningful difference between the policies manifests itself when inflation targets are missed. Then IT results in price-level drift, because targeting a rate does not retain memory of the miss. In contrast, PT remains aimed on the target price-level path.

Specific targeting rules may be derived from general targeting rules. However, the derived rule will change if any of the objective function, constraints, or solution method change. General targeting rules are more flexible because they can react to circumstances. Before turning to general targeting rules, we first examine some specific ones that can be derived from general targeting rules.[8] We classify a rule as a PT rule or an IT rule, according to whether the formula is written in terms of the price level or in terms of inflation.

1.2.1 A specific price-level-targeting rule

Price-stationary rules implement PT policy. Consider a specific PT rule for stabilizing the price level, p_t, around a constant growth target path, $p_t^* = p_{t-1}^* + \pi^*$, where π^* is the implied target inflation rate. The rule can be written as follows:

$$p_t - p_t^* = a(p_{t-1} - p_{t-1}^*) + b\rho u_{t-1} + c\varepsilon_t, \tag{1}$$

where p_t is the logarithm of the price level and *a*, *b*, and *c* are constants. The disturbance is an AR(1) process, $u_t = \rho u_{t-1} + \varepsilon_t$, where $0 \leq \rho < 1$ and ε_t is an i.i.d. shock with zero mean and variance σ^2. The rule allows for different responses to the shock and the lagged disturbance terms.

Whether equation (1) implements PT policy depends on parameters. For $|a| < 1$, $p_t - p_t^*$ is stationary and p_t is price-level-trend stationary. As the price level adjusts towards the target path, the rule implements a PT policy. Series 3 in Figure 1 (section 3) illustrates this case. On the other hand, if $a = 1$, $p_t - p_t^*$ is not stationary, and p_t is not price-level-trend stationary. Then (1) yields $\pi_t = \pi^* + b\rho u_{t-1} + c\varepsilon_t$ so that it is inflation-stationary. As the rule now displays price-level drift, it implements IT policy.

Rule (1) reveals that PT has the effect of "average inflation targeting." The price level and inflation are related:

$$p_t - p_t^* = (t - t_0)[\pi_t^a - \pi^*] + (p_{t_0} - p_{t_0}^*),$$

where

8. Duguay (1994) and Fischer (1994) present and simulate ad hoc targeting rules that are defined only on prices. Svensson (1997c) provides the analytical solution to the variances for these targeting rules. The rules presented in sections 1.2.1, 1.2.2, and 1.2.3 nest these rules.

$$\pi_t^a = \sum_{j=t_0}^{t} \pi_j / (t - t_0)$$

is the average inflation rate and t_0 is the initial period. Rule (1) can be rewritten in terms of average inflation:

$$[\pi_t^a - \pi^*] = a\frac{(t-1-t_0)}{(t-t_0)}[\pi_{t-1}^a - \pi^*] - \frac{(1-a)(p_{t_0} - p_{t_0}^*) + b\rho u_{t-1} + c\varepsilon_t}{(t-t_0)}.$$

Thus, the PT rule implies that average inflation, π_t^a, is stationary around π^* for $|a| < 1$.

Rule (1) can be expressed to examine the dynamics of inflation:

$$\pi_t - \pi^* = -(1-a)(p_{t-1} - p_{t-1}^*) + b\rho u_{t-1} + c\varepsilon_t,$$

where $\pi_t = p_t - p_{t-1}$ is the inflation rate between period $t-1$ and *t*. Thus, if the rule is price-level stationary it is also inflation-stationary around π^*. Suppose that shocks are temporary, $\rho = 0$. When past inflationary shocks move the price, p_{t-1}, above the price-level target, the rule requires disinflation so that current inflation whipsaws below the average π^*. Thus, the rule acts to undo the effect of the shock.

To implement PT, the rule must completely undo the effects of shocks in the long run. How the rule does this is revealed by first differencing equation (1)[9]:

$$\pi_t = a\pi_{t-1} + (1-a)\pi^* + b\rho(u_{t-1} - u_{t-2}) + c(\varepsilon_t - \varepsilon_{t-1}). \qquad (1')$$

It can be shown that the effect of an increment to the shock, $\Delta\varepsilon_t$, on the cumulative change in the price level is:

$$\sum_{j=0}^{\infty} \Delta\pi_{t+j} = c\Delta\varepsilon_t\{1 - (1-a)[1 + a + a^2 + \ldots]\}$$

$$+ \frac{b\rho\Delta\varepsilon_t}{1-a}\{1 - (1-\rho)[1 + \rho + \rho^2 + \ldots]\} = 0.$$

The first term on the right-hand side gives the direct effect of the shock as it affects the economy in successive periods, whereas the second term gives

9. To implement PT, rule (1') must be accompanied with initial values, $p_{t_0-1}^*$ and π_{t_0-1}. Then rule (1') traces the same time path as rule (1) for $t \geq t_0$. If the initial values are consistent with rule (1) at time $t_0 - 1$, then the rule is implemented in a "timeless" way, as if it had always been in effect.

the indirect effect through the lagged persistence term. Both terms sum to zero so that the cumulative change in the price level is zero.

1.2.2 A specific inflation-targeting rule

Consider the following IT rule:

$$\pi_t = a\pi_{t-1} + (1-a)\pi^* + b\rho(u_{t-1} - u_{t-2}) + c\varepsilon_t - d\varepsilon_{t-1}. \qquad (2)$$

This rule is the same as (1') when $c = d$. In this special case, the rule implements PT policy since, as explained above, the effect of a shock is completely undone and there is no price-level drift. However, in the more generic case, $c \neq d$, there *is* price-level drift. Thus, when $c \neq d$ and inflation is stationary, $(|a| < 1)$, rule (2) implements IT by our definition. From this perspective, IT is the more general policy, one that describes all policies that stabilize inflation, except for the one policy that also targets the price level or equivalently, average inflation. Series 2 in Figure 1 illustrates IT for the case $c > d > 0$, where the rule corrects for some of the price-level drift.

It is instructive to examine a simpler version of the rule, where the disturbance does not persist: $\rho = 0$ and $d = 0$. Then, when $a = 0$, the price level is a random walk. This case is illustrated by series 1 in Figure 1 (section 3). When $a = 1$, inflation is a random walk. For $0 < a < 1$, inflation gradually adjusts back to target. Only when $a < 0$ would the rule whipsaw inflation and contain some of the price-level drift. Duguay (1994) and Svensson (1997c) analyze these cases and find that the variance of inflation is less under rule (1) than under rule (2) for some parameterizations. Of course, the price-level variance of the IT rule (2) increases with time, and the unconditional variance is infinite.

1.2.3 Hall's elastic price-target rule

Hall (1984) proposed an "elastic price-target" rule for gradually stabilizing unemployment around the natural rate and the price level around a price-level target. Using Okun's Law, the rule relates the constant price level to the output gap:

$$x_t = -\theta(p_t - p^*), \qquad (3)$$

where the logarithm of the output gap, $x_t \equiv y_t - y_t^n$, is the difference between the log of output and the log of the natural level of output, and $\theta > 0$. This rule is price-level stationary for judicious choice of θ, given a PC. The Bank engages in PT by pushing demand below the natural rate whenever the price level rises above the target. A Bank following such a rule would persist battling past inflation shocks. Series 3 in Figure 1 illustrates such a PT policy response.

Taking the first difference of equation (3) yields:

$$x_t + \theta\ p_t = x_{t-1} + \theta\ p_{t-1}.$$

Thus, the rule stabilizes an index of the price level and the output gap. The "elastic" rule explicitly considers other objectives than just the price level. The rule can be expressed in terms of inflation: $x_t - x_{t-1} = -\theta\pi_t$. Now this is an IT rule. However, the key point is that it implements the same PT policy.

1.3 Target variables and policy procedures

To focus on target rules, the literature we examine usually assumes that the current period target variables—inflation or the price level—are controllable. Thus, some interest rate or money supply instrument rule is assumed to achieve the desired path for the target variables. When this assumption is not satisfied, we follow Svensson (1997a, 1999b) and Svensson and Woodford (1999) who study inflation-forecast targeting. This involves targeting inflation forecasts as the intermediate instrument for policy.

The argument for target rules is that they permit the Bank the flexibility to change instruments when necessary. However, this process may not be transparent and, hence, provides the Bank with the opportunity to deviate from stated policies. In contrast, an instrument rule is easier to monitor. Similarly, general rules permit flexibility but are less transparent than specific rules.

2 Basic Modelling Ingredients

We turn next to the basic ingredients in the Bank's optimization problem, what Svensson and Woodford term "general targeting rules." Such rules use a Bank's objective function, constraints, and an optimization routine to describe how target variables relate to target levels.

2.1 Objective functions

In the literature, the objective functions are usually specified as quadratic loss functions defined on the output gap and inflation. This facilitates work with PC constraints that are specified on the same variables. Trading off inflation and output minimizes losses.

2.1.1 The social loss function

The most commonly used social loss function is

$$\sum_{t=t_0}^{\infty} \beta^{t-t_0} \frac{1}{2}\left[(\pi_t - \pi^*)^2 + \lambda(x_t - x^*)^2\right], \tag{4}$$

where $0 < \beta < 1$ is the discount factor. This function gives the expected value of the sum of discounted future losses starting from period t_0. Society is assumed to value stabilization of both inflation and the output gap, thus the weight λ satisfies $0 < \lambda < \infty$. The targets are exogenous parameters. When $x^* > 0$, output above the natural level of output is desirable. This usually is premised on the existence of inefficiency in the economy.

2.1.2 The bank loss function

It is critical to distinguish the social loss function from the Bank loss function. Following Rogoff (1985), many papers in the literature assume that an objective function can be assigned to the Bank, the authority that formulates and implements specific monetary policy. The assigned function is usually chosen with the aim of minimizing the social loss function. In much of the literature, the Bank aims to minimize a function with inflation as a target variable:

$$E \sum_{t=t_0}^{\infty} \beta^{t-t_0} \frac{1}{2}\left[(\pi_t - \pi^{b*})^2 + \lambda^b (x_t - x^{b*})^2\right], \tag{5}$$

where the superscript *b* identifies parameter values that may differ from the social loss function.[10] Using Svensson's (1997b) terminology, a value of $\pi^{b*} < \pi^*$ indicates an "inflation-target conservative" Bank; a "weight-conservative" Bank is one where $\lambda^b < \lambda$.

A *strict inflation-targeting rule* refers to a Bank that cares only about inflation, i.e., $\lambda^b = 0$. Assuming the target variable π_t can be chosen freely, the Bank can always hit the target, $\pi_t = \pi^{b*}$, and thereby effectively hit the implied target path,

$$p_t = p_t^{b*} \equiv (t - t_0)\pi^{b*} + p_{t_0}.$$

We have our first result for policy.

---

10. Note that the Bank's tenure is infinite, and the discount rate is the same as in the social loss function. These are dimensions that have yet to be explored in the literature.

Result 1. A strict inflation-targeting rule implements a price-level-targeting policy where the price-level path equals the targeted path $\{p_t^{b*}\}_{t=t_0}^{\infty}$.

Is this a good idea? From a social loss perspective, it is clearly inferior under the maintained assumption $\lambda > 0$. When society cares about the output gap, appointing a Bank that also cares about the output gap, if only a little bit, reduces the social loss because it trades off inflation for output stabilization. To examine this trade-off we assume $\lambda^b > 0$, i.e., the Bank pursues a "flexible-targeting" policy.

2.1.3 Target objectives

We say that the Bank has an *inflation-targeting objective* (ITO) if, as above, $\pi_t - \pi^{b*}$ is an argument in the Bank's objective function. It has a *price-level-targeting objective* (PTO) if $p_t - p_t^{b*}$ is an argument in the Bank's objective function. Svensson (1999d), Dittmar et al. (1999), Dittmar and Gavin (2000), Kiley (1998), and Vestin (2000) all use the following Bank loss function with a PTO:

$$E \sum_{t=t_0}^{\infty} \beta^{t-t_0} \frac{1}{2}\left[(p_t - p^{b*})^2 + \lambda^B (x_t - x^{b*})^2\right], \tag{6}$$

where $\lambda^B > 0$ for the same reasons as $\lambda^b > 0$.

Svensson (1999d) refers to equation (5) as inflation targeting and equation (6) as price-level targeting. We use the different language to relate objectives to outcomes. Table 1 anticipates our results, and relates target objectives according to whether the general rule effectively implements IT or PT.

Table 1
Optimal policy by target objective

	Solution	
Bank loss function	**Commitment**	**Discretion**
ITO	IT or PT	IT
PTO	PT	PT

Observe that the PTO always achieves PT. The ITO always achieves an inflation-stationary path, but it may also achieve the more restrictive price-stationary path and thereby implement PT. This occurs if expectations in the PC are strictly forward-looking.

2.2 Phillips curves

The PC details the short-run trade-off between inflation and the output gap. When the Bank is able to choose the target variable, the PC is the only relevant constraint. This is because the Bank is then able to control the entire sequence of x_t and π_t (or equivalently, p_t), the only arguments in the Bank loss function.

2.2.1 A hybrid Phillips curve

Consider a construct that incorporates predetermined and forward elements:

$$\pi_t = \kappa(x_t - \delta x_{t-1}) + \phi E_{t-1}\pi_t + (1-\phi)\beta E_t \pi_{t+1} + u_t, \quad (7)$$

where $\kappa > 0$ is the slope of the PC, $0 \le \delta < 1$ is the coefficient on lagged output, $0 \le \phi \le 1$ is the weight on predetermined expectations, $E_{t-1}\pi_t$, and $E_t\pi_{t+1}$ is the expectation in period t of the inflation rate in period $t+1$. The error term, as before, is an AR(1) process, $u_t = \rho u_{t-1} + \varepsilon_t$, where ε_t is an i.i.d. shock with zero mean and variance σ^2. This construct allows predetermined and forward-looking expectations as well as exogenous and endogenous persistence. These features are critical to whether the transition dynamics to the long-run equilibrium are price-stationary. The sole source of the need for adjustment is the "cost-push" disturbance, u_t.[11]

2.2.2 New-Keynesian and New-Classical Phillips curves

The New-Keynesian PC corresponds to the case when expectations are forward-looking, $\phi = 0$. Expected future inflation, $E_t\pi_{t+1}$, affects the current trade-off between inflation and the output gap. This specification stems from Calvo's (1983) model of staggered nominal-price setting in an environment with imperfect competition. The discount factor, β, enters because suppliers of goods discount future earnings. Firms set prices based on the expectation of future marginal costs. Thus, the disturbance term, u_t, captures anything else that might affect this cost. Using the language of Clarida et al. (1999), we call this a cost-push disturbance. Woodford (1999a) calls this an "inefficient supply shock," because it causes a temporary deviation from the efficient level.

11. Since the PC relates to the output gap, this disturbance excludes shocks to the natural rate (not modelled). For this reason, the disturbance is modelled as non-permanent. Without the cost-push shock, there is no trade-off in the long-run equilibrium. For evidence supporting the presence of the disturbance term and lagged dependent variables in the PC, see Fuhrer and Moore (1995).

The opposite case, $\phi = 1$, is Lucas' (1973) New-Classical PC. In contrast to the above, expected current inflation is predetermined: $E_{t-1}\pi_t$. With predetermined expectations, new developments in period t do not affect the expectations term in the PC at date t and, hence, cannot affect the current PC trade-off through this channel. This specification can be justified from imperfect information, predetermined wages, or a P-bar model (see McCallum 1994).

The different PCs have different implications for monetary neutrality. Taking the unconditional expectation of the hybrid PC with $\delta = 0$ yields:

$$Ex_t = (1-\phi)[E\pi_t - \beta\, E\pi_{t+1}]/\kappa .$$

Money is always neutral in the New-Classical case, since $\phi = 1$ requires $Ex_t = 0$. However, money is not neutral if the New-Keynesian expectations component is given positive weight, $\phi < 1$. For a trend inflation rate of π^*, we get:

$$Ex_t = (1-\phi)(1-\beta)\pi^*/\kappa .$$

Expected output is increasing in the trend inflation rate. McCallum (1994) criticizes this implication of the New-Keynesian PC as unrealistic. Rotemberg and Woodford (1999) admit to being very uncomfortable with the non-neutrality implication and opt for a version of Calvo (1983), where prices are automatically adjusted by the trend inflation in every period. Clarida et al. also use the detrended inflation rate in their analysis. In this paper, we use the original specification, although it is easy to find the results for the detrended specification by simply setting $\pi^* = 0$ and interpreting π_t as the detrended variable. Note, however, that the detrended specification is non-neutral for an increasing inflation rate.

2.2.3 The monetary transmission mechanism

The above description does not detail the role of money and assets because the PC is the only relevant constraint when the Bank can choose the target variables accurately. Any specification of aggregate demand can be used to derive the required setting for the policy instrument. For example, Clarida et al. (1999) use the following IS and LM specification (derived from an optimizing model) in their most general model:

$$x_t = -\sigma(i_t - E_t\pi_{t+1}) + \partial x_{t-1} + (1-\partial)E_t x_{t+1} + g_t$$

$$m_t - p_t = \varphi y_t - \eta i_t + v_t ,$$

where i_t is the nominal interest rate, m_t is log of the stock of high-powered money, g_t and v_t are disturbance terms, and $\sigma > 0$, $0 \le \partial \le 1$, and $\varphi > 0$ and $\eta > 0$ are constants. If the interest rate is the policy instrument, it can be related to the output gap in the IS equation. Alternatively, if money is the instrument, it can be determined using both IS and LM.

The example highlights several points. First, targeting is only possible if the implied instrument settings are feasible. In particular, the zero constraint on the nominal interest rate does not bind. Second, given the Bank's choice of output gap and inflation, the instruments are set to completely offset the money demand shock, v_t, and the IS shock, g_t.[12] Third, accurate targeting is possible only if the Bank knows the functions and disturbances, although if the money shock were unknown, the Bank could use the interest rate as the policy instrument.[13] Finally, since the Bank's problem is separable, nothing in the IS or LM equations—neither forward or lagged components nor shocks—affects the Bank's choice of inflation or output gap. The problem is not separable if the Bank has preferences or constraints for the nominal interest rate (see sections 5 and 6).

Another complication is the existence of lags in the monetary transmission mechanism. It is widely viewed that these lags can be up to two years in duration. In section 3.2 we follow Svensson and Woodford (1999), who model a control lag as arising from predetermined prices. This allows the monetary authority to use inflation forecasts as the optimal intermediate instrument. Again, there is no need to detail the asset and money markets. A fuller analysis of monetary policy would capture the almost immediate price responses in financial and foreign exchange markets and lagged responses in labour and product markets.

2.3 Equilibrium: Commitment vs. discretion

The Bank's problem is to minimize the Bank loss function subject to the appropriate Phillips curve and initial conditions. The two approaches

12. Both Clarida et al. (1999) and Woodford (1999d) derive the "IS equation" from a Euler equation and show that the disturbance term is $g_t = E_t[(y^n_{t+1} - y^n_t) - (q_{t+1} - q_t)]$, where q_t is the shock to autonomous expenditure. Thus, interest rate policy involves completely offsetting temporary shocks to the natural level of output, as well as the more traditional autonomous expenditure. Interest rates move procyclically with consumption or government expenditure shocks and countercyclically to technology shocks. Monetary policy should accommodate shocks to money demand and the natural level of output.

13. With an interest rate rule, nominal indeterminacy is not a problem in our framework, since with sluggish price adjustment the last period's price level is a nominal anchor. Also, the implied interest rate rule provides feedback from endogenous variables (inflation and the output gap), a condition for determinacy. See Clarida et al. for a discussion of real indeterminacy.

commonly used to solve this problem lead to quite different outcomes. The open-loop solution involves the Bank choosing in some initial period, say t_0, a state-contingent plan for the control variables,

$$\{\pi_t, x_t\}_{t = t_0}^{\infty}.$$

This is often referred to as the "commitment" solution, because the Bank sticks with the plan, come what may, in all subsequent periods. In contrast, the "discretion" or closed-loop solution involves the Bank re-optimizing in each period by choosing a new state-contingent plan consistent with the new initial conditions.

In a classic paper, Kydland and Prescott (1977) show that with rational expectations a policy-maker with discretion has the incentive to renege on the previous optimal plan and to adopt a new one. Because the private sector can anticipate this, plans lack credibility and cannot be used to beneficially influence expectations. The equilibrium has a Prisoner's Dilemma quality. In contrast, commitment or rules can be used to influence future expectations leading to higher welfare. The advantage of discretion is that it allows the Bank to respond to unforeseen changes in the economy. However, this advantage cannot be captured in the model, since the structure of the model is assumed to be fully known.

There is considerable debate over whether commitment or discretion is the better description of the Bank's problem. McCallum (1995, 1996) argues that central banks can simply commit to policy. On the other hand, Svensson (1999d) counters that this reasoning does not generate subgame-perfect equilibria in standard models. Chari et al. (1989) go further and argue that without a commitment technology, discretionary policy is the *only* behaviour of the policy-maker consistent with a rational-expectations equilibrium. An intermediate view would seem reasonable, but without more institutional structure, serious criticisms can be levelled against all approaches.[14]

An intermediate approach, which is prominent in the literature, assumes that the Bank must act under discretion, but can be assigned a Bank loss function that mitigates against the suboptimality of discretion. A leading example is Rogoff's (1985) famous argument for appointing a weight-conservative Bank, $\lambda^b < \lambda$. Svensson (1999d) and others take up this approach in the PT literature. They argue for assigning the Bank a PTO loss function like equation (6). This approach suffers from schizophrenia in that

14. Howitt (2000) provides a novel discussion of learning and games theory applied to monetary institutions.

it implies a commitment technology in the assignment of policy but not in its execution.

Woodford (1999a, 1999c) takes an interesting stance that leans towards commitment but possesses the advantages of discretion. He argues for using specific rules but changing them in ways that provide no incentive to deviate from stated policy objectives.[15] Here, transparency and experience with the policy process determine its credibility. This approach may suffer from leaving too much leeway in the assignment of policy.

The debate over rules versus discretion has considerable practical implications for how a country should design monetary institutions for credibility and flexibility. In our analysis, we emphasize the commitment approach, because we believe it is the less problematic of the two approaches. However, we also examine the discretion literature and indicate when both approaches recommend similar policies.

3 Commitment: Should the Bank Stabilize the Price Level?

This section describes a Bank that can credibly commit to a policy for the indefinite future. As a prerequisite, of course, the Bank should be assigned the social loss function (4). Thus, the Bank has an ITO.

3.1 Targeting current variables

3.1.1 New-Classical vs. New-Keynesian Phillips curves

Svensson (1999d) shows that for a New-Classical PC ($\phi = 1$ and $\rho = 0$), the optimal commitment policy involves IT. Clarida et al. (1999) and Vestin (2000), in contrast, show that the optimal commitment policy involves PT for a New-Keynesian PC ($\phi = 0$ and $\delta = 0$). The key to the different outcomes stems from the expectations specifications of the models.

The New-Keynesian PC has forward-looking expectations, $E_t\pi_{t+1}$, which are affected by the Bank's future policy actions and feedback to influence the current trade-off between inflation and output. The Bank can improve this trade-off by committing to PT. Consider, for example, an

15. Woodford (1999a, 1999c) advocates "timeless" rules that do not depend on initial conditions (e.g., specific rules (1)–(3)). Changing rules midstream then involves implementing the new rule as if it had been in place all along. Because such rules do not have an initial condition they eliminate the incentive to opportunistically deviate. However, such rules may saddle society with past burdens. A Bank that can signal that it would not deviate, would do better with a rule that considers current initial conditions.

inflation shock, $\varepsilon_t > 0$. A credible commitment to eventually bring the price level back to the target path lowers expectations of inflation, including the current forecast, $E_t\pi_{t+1}$. This, in turn, lowers the current inflation rate in equation (7).

By contrast, if the inflation-output trade-off depends on predetermined expectations, $E_{t-1}\pi_t$, as in the New-Classical PC, the policy response to current shocks does not feed back through expectations on the current PC trade-off. The optimal commitment policy is less stringent and involves IT. Quite simply, there is nothing to be gained (in terms of today's inflation-output trade-off) by a hard-nosed, PT position.

3.1.2 The hybrid commitment solution

What effects dominate if the expectations terms $E_{t-1}\pi_t$ and $E_t\pi_{t+1}$ are weighted as in the hybrid PC? We will outline the Bank's problem and show that unless expectations are completely forward-looking, IT is optimal. To simplify the analysis, the endogenous persistence term $(\delta = 0)$ is omitted from the PC.

The Bank's problem is to choose the sequence

$$\{\pi_t, x_t\}_{t=t_0}^{\infty}$$

to minimize the social loss function equation (4) subject to the hybrid PC equation (7):

$$E_{t_0}\sum_{t=t_0}^{\infty}\beta^{t-t_0}\frac{1}{2}\left[(\pi_t-\pi^*)^2+\lambda(x_t-x^*)^2\right]$$

$$\text{s.t. } \pi_t = \kappa x_t + \phi E_{t-1}\pi_t + (1-\phi)\beta E_t\pi_{t+1} + u_t .$$

The first-order conditions are:

$$\lambda(x_t - x^*) - \kappa\psi_t \le 0, \tag{8}$$

$$\pi_t - \pi^* + (\psi_t - \phi E_{t-1}\psi_t) - (1-\phi)\psi_{t-1} \le 0, \tag{9}$$

where ψ_t is the Lagrange multiplier at date *t*. Since there is no constraint in period $t_0 - 1$, we have $\psi_{t_{0-1}} = 0$. This initial condition indicates that the Bank ignores the history prior to the initial period, and in the initial period acts as it would in the discretion solution.

The optimal inflation rule is derived from conditions (8) and (9),

$$\pi_t = \pi^* - \frac{\lambda}{\kappa}\{[x_t - \phi E_{t-1}x_t] - (1-\phi)x_{t-1}\}. \quad (10)$$

The optimal inflation policy uses a feedback rule from the current and lagged (expected and actual) output gap. In the New-Keynesian case $(\phi = 0)$, the feedback rule is

$$\pi_t = \pi^* - (\lambda/\kappa)\{x_t - x_{t-1}\},$$

which is just a version of Hall's rule for $\theta = (\kappa/\lambda)$.[16]

In the New-Classical case $(\phi = 1)$, the PC and equation (10) yield

$$\pi_t = \pi^* + \frac{(\lambda/\kappa^2)}{1 + (\lambda/\kappa^2)}\varepsilon_t.$$

For later comparison, we write Svensson's (1999d) solution, which has endogenous persistence $(\delta \geq 0)$, but not exogenous persistence $(\rho = 0)$:

$$\pi_t = \pi^* + \frac{(\lambda/\kappa^2)}{1 + (\lambda/\kappa^2) - \beta\delta^2}\varepsilon_t. \quad (11)$$

Both solutions take the form of the IT rule (2) with $a = b = d = 0$. The price level is a random walk. The presence of endogenous persistence, $\delta > 0$, makes the Bank's job more difficult and induces more price-level drift. However, even as $\delta \to 1$, the Bank maintains the variance of inflation finite. We later contrast this with the case of discretion, where the inflation variance becomes infinite.

We now consider the hybrid case, $\phi < 1$. The Appendix details the derivation of the commitment solution. Given initial conditions, the optimal inflation path can be described by an IT rule (2):

$$\pi_t = a\pi_{t-1} + (1-a)\pi^* + b\rho(u_{t-1} - u_{t-2}) + c\varepsilon_t - d\varepsilon_{t-1},$$

where $0 < a < 1$, $b > 0$, $c > 0$, and $d > 0$. As $a < 1$, inflation is stationary. The initial shock has a positive impact on inflation, $c > 0$. The Bank responds in the next period by acting to correct a portion of the shock as $d > 0$.[17]

16. See section 1.2 and Woodford (1999a) for a discussion of Hall's rule. Note that this version of the rule is consistent with an AR(1) disturbance and $\pi^* > 0$.

17. The model treats the components of the shock to the PC differently. Predetermined expectations are conditioned on the forecastable portion of the AR(1) process, ρu_{t-1}, while forward-looking expectations are conditional on u_t. Hence, the components appear separate in this equation.

Only in the strict New-Keynesian case $(\phi = 0)$ is the price level stationary. Then $b = c = d$ and equation (2) takes the form of equation (1'). The initial response to the shock is eventually fully corrected, $c = d$. Also, the response to the shock and the autocorrelated disturbance is undifferentiated, $b = c$. In the generic case, $\phi > 0$, the disturbance term is differentiated, and $c > d$, so there is price-level drift. The amount of the drift diminishes as $\phi \to 0$. Thus, IT is generally optimal, but it takes the form of correcting some portion of the impact of the shock in subsequent periods.

Figure 1 illustrates the two polar cases and the hybrid case for a 1 per cent inflation shock in period 1. Figure 2 illustrates how the coefficients *a*, *c*, and *d* behave as a function of ϕ. The parameters used are: $\beta = 0.95$, $\lambda = 1$, $\kappa = 0.3$, $\rho = 0$, and $\pi^* = 0$.[18] In the New-Classical case (series 1, Figure 1), the Bank mutes the inflation effect of the shock in period 1, but then ignores the effect of the shock in subsequent periods. This results in price-level drift, but output is restored to the natural rate in period 2. In contrast, in the New-Keynesian case (series 3), the Bank mutes more of the inflation and output shock in period 1. In subsequent periods, the Bank whipsaws inflation below trend. Since the trend inflation is zero, a deflationary policy is maintained until the original price level, $p^* = 0$, is restored. This deflationary policy maintains output below the natural rate in the transition. The hybrid case yields intermediate results and some price-level drift.

Result 2. Inflation targeting is optimal when predetermined expectations are present, $0 < \phi \leq 1$. To the extent that expectations are forward-looking, the optimal policy involves undoing some of the price-level impact of the shock in subsequent periods to the shock $(d > 0)$. However, only in the New-Keynesian case $(\phi = 0)$, where expectations are strictly forward-looking, is it optimal to undo the entire impact of the shock ($c = d$), so that PT is optimal. At the other polar extreme, the New-Classical case $(\phi = 1)$, none of the impact of the shock is corrected in subsequent periods ($d = 0$).[19]

What about optimal output? Equation (A2) in the Appendix details the solution for optimal output. The commitment solution does not contain

18. Using one year as the measure of time, we choose $\beta = 0.95$. Following Williams (1999) and others, we weight the objectives equally, $\lambda = 1$. We choose $\kappa = 0.3$ as Roberts (1995) estimates $\kappa \in$ (0.2, 0.4). Finally, we set $\pi^* = 0$ to make the figure comparable with the case where π_t is interpreted as detrended inflation.

19. Unfortunately, we cannot provide a generic restatement of Result 2, allowing for a lagged output term in the PC, since an analytical solution to the Bank's problem cannot be obtained for $0 < \delta < 1$. We conjecture that endogenous persistence would not change the above result towards PT—although it induces forward-looking behaviour. Only in the New-Classical case do we get a solution with $0 \leq \delta < 1$ in equation (11). Here, lagged output has the effect of increasing price-level drift.

Figure 1
Price-level, inflation rate, and output responses to a supply shock
Series 1 = New-Classical; Series 2 = hybrid; Series 3 = New-Keynesian

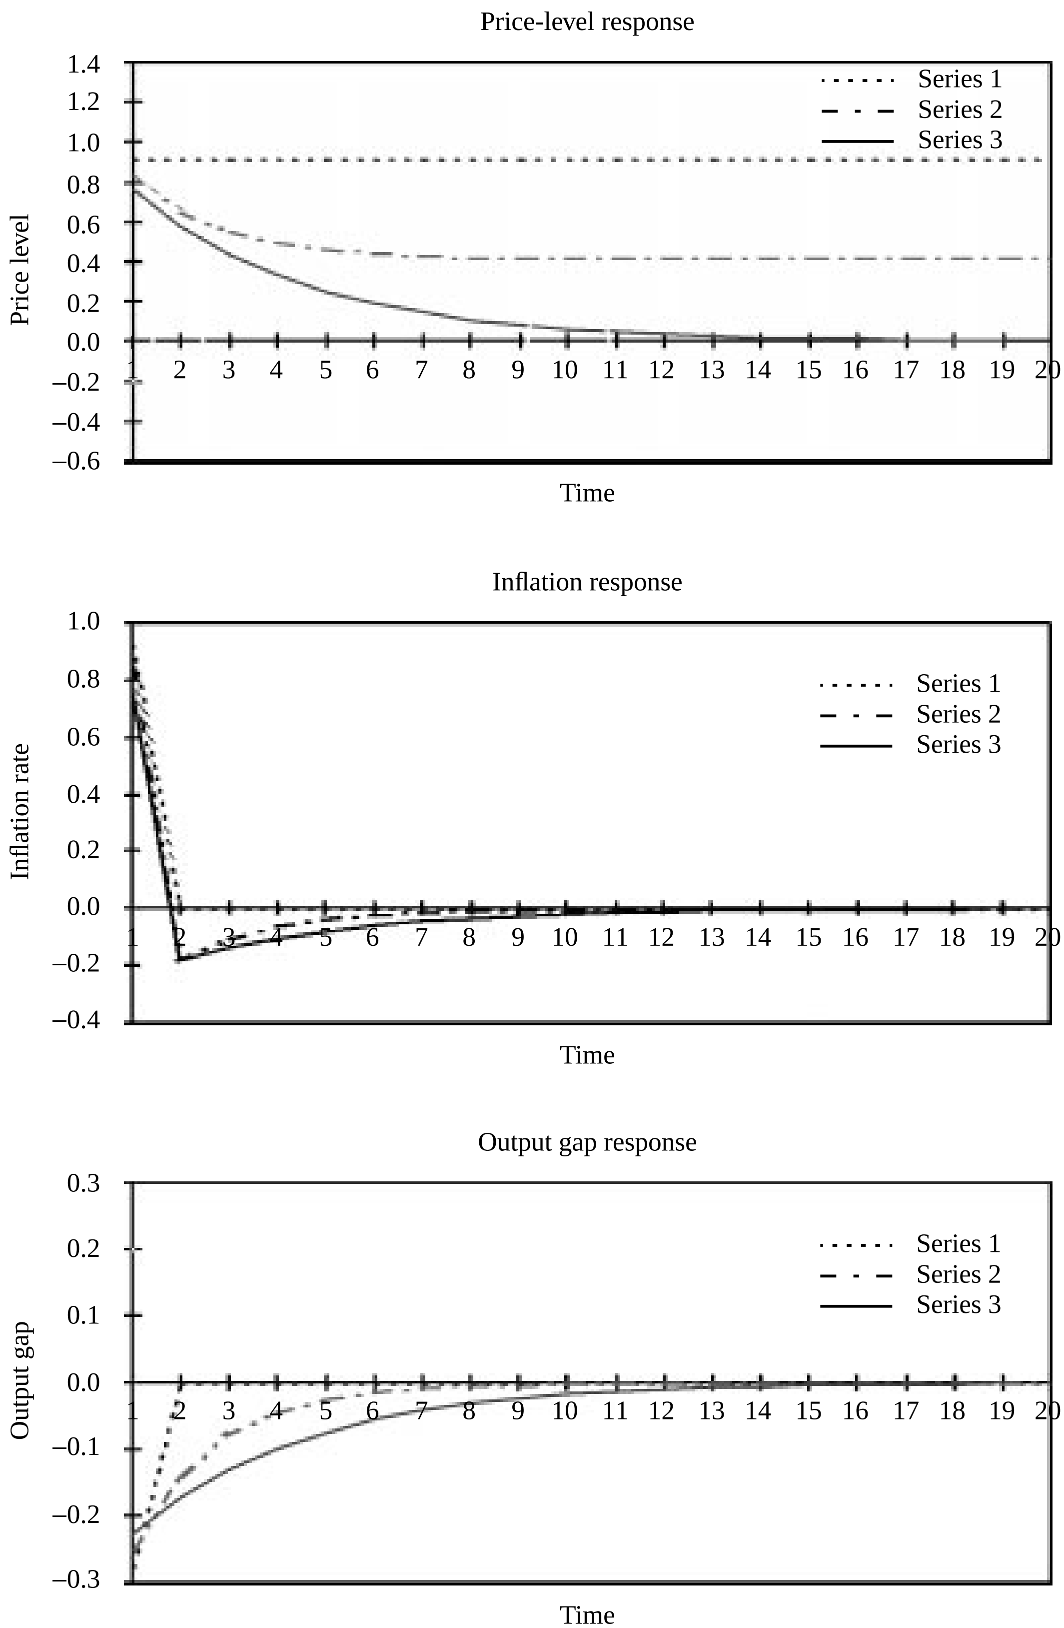

Figure 2
Coefficients *a*, *c*, and *d* as a function of ϕ
(the weight on predetermined expectations)
Series 1 = *c*; Series 2 = *d*

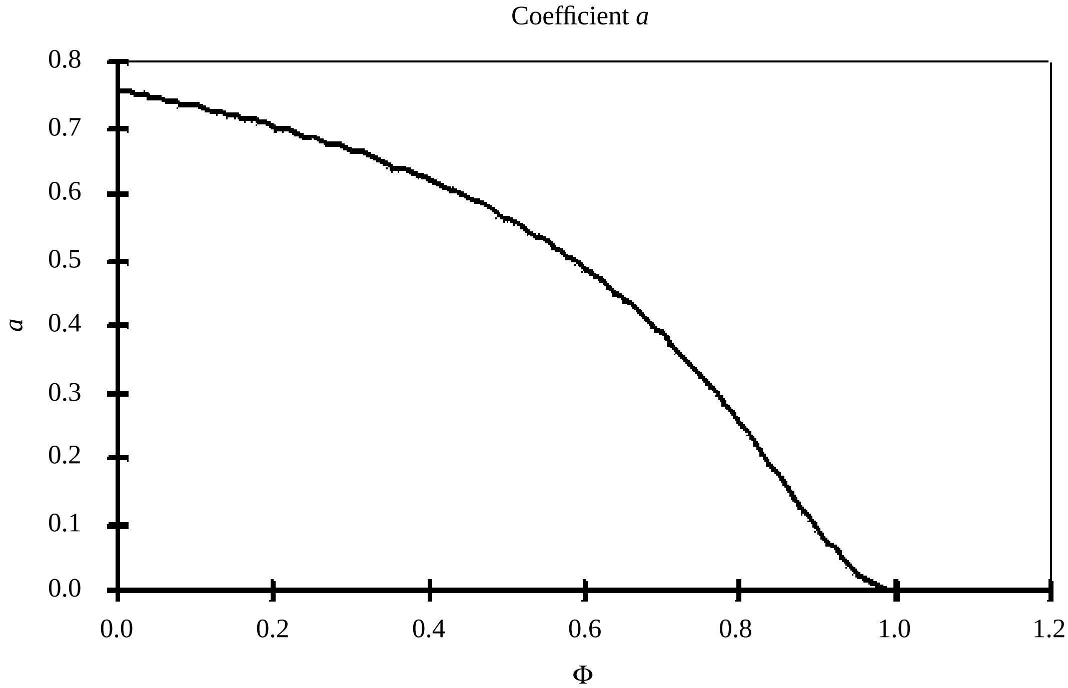

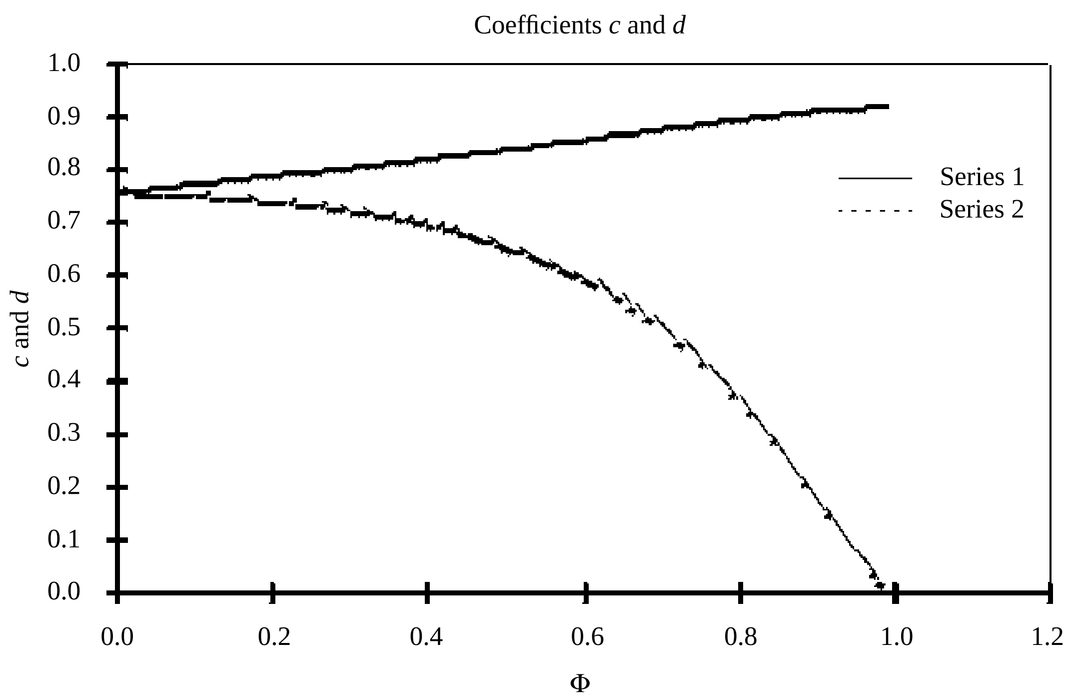

the target output level, x^*. However, it does contain π^*, and is non-neutral unless $\phi = 1$ or $\beta = 1$. As explained in Section 2.2, the forward-looking-expectations component in the hybrid PC forces the following non-neutral long-run relationship: $Ex_t = (1-\phi)(1-\beta)\pi^*/\kappa$. Most economists find the non-neutrality of money in the long run unrealistic, and this non-neutrality can be removed from the model by treating π_t as detrended inflation and setting $\pi^* = 0$.

In our optimization problem, π^* was taken as exogenous. But when using the long-run relationship, it is optimal to choose the inflation target to hit the output target: $\pi^* = \kappa x^*/(1-\phi)(1-\beta)$. Then the Bank can hit both inflation and output targets in the long run. When the economy is efficient such that $x^* = 0$, the optimal inflation target is $\pi^* = 0$, which is consistent with Woodford (1999c), as discussed in section 5.1.

3.1.3 Commitment to the past as well as to the future . . .

The optimal commitment solution involves sticking to the state-contingent plan set in the initial period. Woodford (2000) stresses that optimal policy is history-dependent. To understand this, note that the Hall rule can be written in terms of the control variables as $x_t = -\theta(p_{t-1} + \pi_t - p^*)$ so that the lagged price enters the rule. The New-Keynesian adjustment path in Figure 1 follows Hall's rule—the Bank persists in maintaining a negative output gap as long as the price level is above the target. In the opposite, New-Classical case $(\phi = 1)$, optimal policy is history-dependent when there is endogenous persistence from lagged output in the PC. This is because expectations formed in the current period depend on the current output choice.[20]

Result 3. The optimal commitment policy is generally history-dependent.

What is the nature of the optimal commitment to the future? History-dependent policy often implies that the optimal policy promises a gradual adjustment of output and inflation. This can be seen in Figure 1, where the period 1 output and inflation responses to the shock are muted relative to the New-Classical case. The credible commitment to deflate in subsequent periods lowers inflation expectations in period 1. Interestingly, Clarida et al. (1999) show that interest rates also adjust gradually in the initial period, by

20. The New-Classical case without persistence is not history-dependent, but this is not an important exception, because, as Dittmar et al. (1999) show, the empirical evidence supports substantial persistence.

less than the inflation shock.[21] Woodford (1999a) shows that the persistent policy leads to interest rate inertia.

3.2 Inflation-forecast targeting when prices are predetermined

When the monetary transmission mechanism has substantial lags, the Bank's problem involves anticipating the control lag for policy and using forecasts to guide the economy.[22] If the control lag is one period, the Bank should use $E_{t-1}\pi_t$ as its target variable in period $t-1$. Following Rotemberg and Woodford (1997, 1999) and Svensson and Woodford (1999), the control lag arises from the fact that prices are predetermined in the current period. In their model, this implies that current inflation and output are predetermined. Furthermore, there is no feedback from future policy to current variables through aggregate demand. With such a structure, Svensson and Woodford show that minimizing the loss function involves choosing the inflation and output forecasts to minimize

$$E_{t_0}\sum_{t=t_0}^{\infty} B^{t+1-t_0}\frac{1}{2}\left[(E_{t-1}\pi_t)^2+\lambda(E_{t-1}x_t-x^*)^2\right].$$

To simplify the analysis, we assume that $x^* = 0$. We find that when the Bank cannot react to the current shock in a way that affects current variables, it is optimal to absorb the shock into the price level.

Result 4. With predetermined prices and output, PT is not optimal.

We only consider the New-Keynesian case $(\phi = 0)$, since this is the only one where PT is optimal without predetermined prices. Svensson and Woodford (1999) write the modified PC consistent with predetermined prices as

$$\pi_t = \kappa E_{t-1}x_t + \beta E_{t-1}\pi_{t+1} + u_t.$$

21. Since this pattern looks like the discretion response, the authors contend that it would not be credible. Instead, they argue for a restricted optimal policy where the interest rate rises by more than the shock in the initial period.

22. In a series of papers, Svensson (1997a, 1999a, 1999b, 1999c) makes a case for "inflation-forecast targeting." This policy refers to using forecasts of the target variable—inflation—as an intermediate target variable. He argues that the conditional inflation forecast corresponding to the control lag is the best intermediate instrument for picking future inflation. Svensson (1997a, 1999b) examines a model where policy affects output with a one-period lag and inflation with a two-period lag. This captures the common belief that it takes twice as long for policy to affect inflation as it does to affect output. Dehejia and Rowe (2000) examine such a lag structure and find that output variability is lower when targeting the forecast price level rather than forecast inflation. Here, we examine a situation with a simpler control lag of one period for inflation as well as output.

Current variables can only be affected by choices made in the previous period.

Inflation-forecast targeting involves choosing $E_{t-1}\pi_t$. Targeting is imperfect because of the current period shock, ε_t, so that actual inflation differs from the target according to $\pi_t = E_{t-1}\pi_t + \varepsilon_t$. The analysis is straightforward when there is no persistence. Then

$$E_{t-1}\pi_t = \kappa E_{t-1}x_t + \beta E_{t-1}\pi_{t+1}.$$

To minimize the loss, set $E_{t-1}\pi_t = E_{t-1}\pi_{t+1} = 0$, to achieve $E_{t-1}x_t = 0$. Thus, $\pi_t = \varepsilon_t$ and there is price-level drift, $p_t = p_{t-1} + \varepsilon_t$; the shock is sunk and not worth undoing.

With persistence in the disturbance, the predictable part of the error has an impact that the Bank can ameliorate by trading off inflation and output. This predictable part can be targeted just like in the problem where current variables can be targeted. The only difference is that the disturbance the Bank conditions on is ρu_{t-1}, instead of u_t, and the Bank target variable is $E_{t-1}\pi_t$, instead of π_t. Hence, the predictable error is price-stationary, but the overall shock results in drift, $lim_{j\to\infty} E_t p_{t+j} = p_{t-1} + \varepsilon_t$.

4 Discretion: Free Lunch?

Discretion characterizes the situation when the Bank cannot credibly commit to a policy for the indefinite future. It is well-known that this situation leads to an *inflation bias*, when the Bank targets output above the natural rate, $x^* > 0$. Such a Bank has an incentive to engineer inflation surprises to increase output. But the public is not fooled. The inefficient equilibrium displays the natural rate of output, $x = 0$, with a level of inflation above the target, $\pi > \pi^*$.

Less well-known is that a problem remains even when the Bank targets the natural rate, $x^* = 0$. The equilibrium displays an inefficient response to shocks, what Svensson (1997b) refers to as a *stabilization bias*. This arises because the social loss function is defined on an ITO. With this objective, the Bank is unable to condition future expectations in a desirable way, because its dynamic program ignores history.

Proposed remedies for these biases take the form of appointing a Bank with a different objective function. The famous remedy for the inflation bias proposed by Rogoff (1985) is to appoint a weight-conservative Bank with $\lambda^b < \lambda$. Svensson (1997b) explores a number of remedies for the stabilization bias by modifying an ITO Bank loss function. He finds that a weight-conservative Bank can also effectively reduce the stabilization bias.

Although modifying an ITO Bank loss function helps, the Bank's dynamic program still ignores history under discretion.

This last point suggests that there may be some merit in modifying the objectives of the Bank loss function to track history. Svensson (1999d) provides the interesting result that assigning the Bank a PTO loss function, equation (6), improves social welfare by reducing the inflation variation without worsening output variation. He refers to this as a "free lunch," although we use the term to refer to a welfare improvement with a PTO. The free-lunch result has been developed by a number of authors and is related to the role of price stationarity in reducing the variability of inflation, which ameliorates inflationary and stabilization bias under discretion.

4.1 Free lunch with New Classicalists and New Keynesians

Svensson's (1999d) original analysis demonstrates the free-lunch result in a very specific situation: for a New-Classical PC with endogenous output persistence, $\delta > 0.5$, and for preference weights on output that are the same across ITO and PTO loss functions, $\lambda^B = \lambda^b$. Dittmar et al. (1999) and Parkin (2001) generalize the analysis to show that for any choice of λ^b there exists a λ^B that yields both lower output and inflation variance, provided that output persistence is sufficiently high. Dittmar and Gavin (2000) and Vestin (2000) provide such a general analysis in the New-Keynesian case, and show that endogenous persistence in the PC is not always needed for the free-lunch result. In addition, the PTO loss function not only has the advantage of reducing the stabilization bias when $x^* = 0$, but it also eliminates the inflation bias when $x^* > 0$.

The free-lunch result obtains when using a PTO loss function better mimics the commitment solution than does using an ITO loss function. This intuition is clearest in the New-Keynesian case, where the commitment solution is price-stationary, but the discretion solution is not. Using a PTO loss function under discretion is preferable because it is also price-stationary. To see this analytically, recall that the optimal inflation rule (10) implies

$$\pi_t - \pi^* = (\lambda / \kappa)(x_t - x_{t-1}).$$

By contrast, with discretion, the rule is $\pi_t - \pi^* = (\lambda / \kappa) x_t$.[23] The commitment rule incorporates history because it tracks the change in the output gap, $x_t - x_{t-1}$, whereas the discretion rule ignores the history by

23. Under discretion, the Bank takes private sector expectations as given. The first-order condition for inflation in section 3.1 becomes $\pi_t - \pi^* + \psi_t \le 0$. Combining this with equation (8) yields the optimal inflation rule.

only considering the level of the output gap, x_t. Under discretion, using the PTO loss function (6) yields an inflation rule which, like the commitment rule, is defined on the change in the output gap, $x_t - x_{t-1}$.[24] For this reason, the PTO loss function better mimics the commitment solution. In fact, Vestin (2000) shows that in the special case of no persistence, a λ^B can be found that replicates the commitment solution. Otherwise, the PTO may mimic, but be unable to replicate, the commitment solution.

The same reasoning applies in the New-Classical case with the proviso that there is enough endogenous persistence. The commitment solution in this case is a function of $x_t - \delta x_{t-1}$. When δ becomes large, this term approximates the change in the output gap, $x_t - x_{t-1}$. Then the PTO does a good job of mimicking the commitment solution. A PTO does better than an ITO for the same reason: an ITO yields an inflation rule that depends on the level of the output gap, while the rule under discretion depends on a change in the gap.

The free-lunch result obtains because PT acts to condition expectations to prevent inflation and output variability from rising. This works in the New-Classical case even though expectations are predetermined, because rational expectations are indirectly forward-looking. The choice of output today affects the next period's Phillips curve through the lagged output term. The lagged persistence must be endogenous. With exogenous output persistence, the Bank can do nothing to affect the inflation-output trade-off.

When the free-lunch result obtains in the New-Classical case, it does so even though the optimal commitment policy involves price-level drift. The result has more force in the New-Keynesian case, since expectations are directly forward-looking, and the commitment solution does not involve price-level drift. Hence, persistence in the PC may not be needed for the result. However, Dittmar and Gavin's (2000) simulations demonstrate that endogenous persistence improves the performance of the PTO over the ITO. The evidence strongly suggests the following result.

Result 5. A PTO under discretion can improve upon an ITO under discretion, provided the degree of persistence in the output gap is large enough and the Bank is assigned the appropriate weight, λ^B, in the loss function.

From another perspective, the free-lunch result holds when the discretion solution with an ITO does badly. Consider the New-Classical

24. First differencing Vestin's (2000) optimal price-level rule in the case of no persistence yields $\pi_t - \pi^* = C(x_t - x_{t-1})$, where C is a function of the parameters. Vestin shows that it is possible to find a λ^B that yields the commitment solution, which amounts to $C = \lambda/\kappa$.

case. Here, inspection of Svensson's (1999d) results reveals that the commitment solution is not price-stationary; nevertheless, a PTO may improve matters. As δ approaches 1, the commitment solution (11) and the PTO solution yield finite inflation variances. In contrast, the ITO under discretion yields an infinite inflation variance. Although the PTO may poorly mimic the commitment solution, it does so much better than the ITO.

4.2 Free lunch, anyone?

Under discretion, the possibility of assigning the Bank a loss function with a PTO suggests that there are other rules that deliver a free lunch. Woodford (1999d) shows that introducing a nominal-interest-smoothing term, $i_t - i_{t-1}$, into the loss function alleviates the suboptimality of discretion by generating greater interest rate inertia. Jensen (1999) provides another example of an alternative discretionary general rule that yields the free-lunch result. He shows that a Bank that stabilizes the growth rate of nominal income may do better than a Bank that attempts to minimize the social loss function under discretion. Introducing the growth rate of nominal income into the Bank loss function forces the Bank to care about gross real output growth. This introduces history into the Bank loss function and yields a Bank that persists in dampening inflation shocks.

Svensson (1997c) mentions that stabilizing the level of nominal income has similar properties to PT. Consider such a Bank using the following period loss function:

$$[Y_t - Y_t^*]^2 = [(p_t - p_t^*) + (x_t - x^*)]^2 = [(p_t - p_t^*)^2 + (x_t - x^*)^2] + 2(p_t - p_t^*)(x_t - x^*),$$

where income Y_t and target Y_t^* can be decomposed into the log price and log output. Thus, targeting the level of nominal income reduces to the PTO function (6) when $\lambda = 1$, except for the cross term. From this perspective, targeting the level of nominal income appears to be a special case of a Bank with a PTO. However, it seems to have the advantage of transparency. Hall and Mankiw (1994) discuss other merits of nominal income targeting.

The possibility of assigning the Bank a loss function dramatically different from the social loss function suggests even better rules. Why not assign the Bank Hall's rule (3)? It is simple and transparent and implements the commitment solution in the New-Keynesian case. In the hybrid case, it would be suboptimal, but would dominate a PTO because it precludes discretion. Better yet, assigning the Bank rule (2) or (10) implements the commitment solution for the hybrid model.

Of course, the argument could be made that such rules do not permit flexibility. But presumably the benevolent authority that mandated the rule in the first place could dictate changes as needed. This perhaps introduces discretion through the back door, but reveals the tension in the logic: assigning general rules with discretion is okay, but not specific rules. Outside of the optimizing framework, discretion is argued to be valuable for reacting to unforeseeable situations. But then the tricky case has to be made that a Bank with a loss function different from the true social loss function would do at least as well in reacting to these unforeseeable situations.

5 Would the Real Social Loss Function Please Stand Up?

The analysis to this point rests on the assumption that the social loss function has an ITO and takes the form of equation (4). The main reason for using this function seems to be that others use it as a tractable method of weighing policy objectives.[25] Recently, the quadratic loss function has been put on a much stronger microeconomic footing. Nevertheless, the analysis seems to miss important elements that suggest a PTO might be included in the social loss function.

5.1 Microfoundations

5.1.1 The New-Keynesian case

Recently, Rotemberg and Woodford (1997, 1999) and Woodford (1999c) have derived the social loss function (4) from an underlying optimizing model. The model has an infinitely-lived representative consumer-worker with period utility defined on Dixit-Stiglitz preferences over differentiated goods. Production is given by Calvo's (1983) model of monopolistic competition and staggered-price setting. The general equilibrium has consumption equalling production in each period, since there is no capital. The natural welfare criterion is expected utility.

Taking a quadratic approximation of expected utility yields the social loss function (4), where $\pi^* = 0$, $x^* = 0$, $\lambda = \kappa / \theta$, and θ is the elasticity of demand facing the average producer. Thus, the analysis not only gives us the form of the loss function but the long-run inflation target and the relative weight on output. The concerns of policy should be to stabilize the output gap and inflation, treating expected and unexpected inflation equally. The

25. Clarida et al. (1999, 1669) write: "Judging by the number of papers written by Federal Reserve economists that follow this lead [quadratic loss function], this formulation does not seem out of sync with the way monetary policy operates in practice (at least implicitly)."

long-run target for output is zero due to taxes eliminating the inefficiency of production. Rotemberg and Woodford (1997) estimate $\lambda = 0.05$, when both the output gap and inflation are measured in percentage points. Thus, they find distortions associated with inflation are much more important.

Inflation enters the loss function because it is the cause of price dispersion, which in turn affects output levels. In an otherwise efficient economy, the long-run inflation target should be zero, because then there is no dispersion arising from staggered-price setting. Woodford (1999c) examines the case where there is no cost-push supply shock (although its presence does not affect the form of the loss function). In this case, optimal commitment policy involves PT, where the Bank sets inflation equal to zero in all periods.[26] However, this result depends on disturbances having a symmetric effect on all sectors and all sectors having the same degree of price stickiness. This caveat would almost surely apply to our analysis when the cost-push shock is included.

5.1.2 The New-Classical case

Woodford (1999c) also examines a New-Classical model where a portion of firms set prices one period in advance and the others flexibly adjust prices. This yields a New-Classical PC. The corresponding quadratic loss function is the same as above except for the inflation target, $\pi_t^* = E_{t-1}\pi_t$. Now, the policy objective should aim to minimize the inflation surprise, $\pi_t - E_{t-1}\pi_t$, in every period. Woodford examines the case without a cost-push supply shock, u_t (although its presence does not affect the form of the loss function). In this case, there is no conflict between objectives and setting inflation to whatever was expected in the initial period, and zero thereafter minimizes the social loss. Thus, the optimal policy is essentially PT.

Result 6. In Woodford's (1999c) New-Classical and New-Keynesian models, targeting a fixed price level is optimal in the absence of a cost-push shock. This is because the Bank is able to eliminate the dispersion of prices by targeting zero inflation.

In the standard framework without a cost-push shock, PT is also optimal (ignoring initial conditions). This is because there is no short-run trade-off between inflation and output variability. With a cost-push supply shock, we know that PT is optimal in the New-Keynesian case. Consider Woodford's New-Classical model with a cost-push shock. Then PT is not

26. King and Wolman (1999) also analyze a utility-maximizing model with staggered-price setting and find that targeting a fixed price level is the optimal policy under commitment. The shock in their model is a productivity shock that affects marginal costs directly and, therefore, is not an inefficient supply shock.

optimal, since expectations are set in advance so that the Bank cannot condition expectations to the current shock.

5.1.3 Nominal-interest-targeting objective

Woodford (1999c) extends his New-Keynesian model to include real balances as an argument in the period utility function. Real balances proxy for the value of liquidity services provided by money. The derived social loss function now has a nominal interest rate objective around a long-run nominal target, $i^* < 0$, the level that minimizes the transaction costs of holding money. Of course, this is not feasible because of the zero nominal interest rate floor, $i \geq 0$.

Woodford (1999c, 1999d) sets $i^* > 0$ sufficiently large that the non-negativity condition does not bind. There is now a conflict between raising the interest rate to fight inflation and stabilizing the interest rate. Interestingly, optimal policy handles this conflict by overcorrecting an inflationary shock such that there is an eventual decline in the price level. This further lowers inflation expectations and, hence, the inflation impact of the shock. The policy also acts to reduce the impact on the nominal interest rate, since the nominal interest rate incorporates the inflation rate.

Overcorrecting the shock is optimal because reducing interest rate variability is another reason to condition expectations for lower future inflation. Here, PT is not optimal because it does not overcorrect the shock. But insofar as PT fully corrects for the shock, it is substantially in the right direction.[27] An interesting exercise would be to modify our hybrid PC model to determine the balance of weights on forward expectations and the nominal interest rate objective that would yield PT as optimal.

5.2 A price-level-targeting objective in the social loss function

5.2.1 Implications of a PTO

The case for PT would be very strong if the true social loss function contained a PTO, $p_t - p_t^*$. A Bank that is assigned the PTO loss function (6)

27. Woodford and Rotemberg (1999) use the standard loss function but restrict interest rate variability. They find that the price-level correction may involve an eventual price decline that is twice as large as the original inflationary shock!

always engages in PT.[28] If a PTO were added as an additional argument into the standard Bank loss function (5), PT would also be the optimal outcome. The reason is that in the long run there is no optimal trade-off between inflation, output, and the price level. With a quadratic loss function, moving towards the target price-level path would always reduce the social loss in the long run. However, also adding a nominal interest rate objective upsets this result, since this objective supports a policy that overcorrects shocks (see previous section).

Suppose that the social loss function took the form of equation (6). Would assigning the Bank this social loss function under discretion be optimal? Vestin (2000) shows that it would generally not be so because the Bank could not completely condition expectations as desired.

5.2.2 Rationale for a PTO

Rotemberg and Woodford (1997, 1999) and Woodford (1999c) find that the social loss function contains an ITO to reduce the variability of price dispersion. With staggered-price setting, a proportion of firms change their price in every period independent of the duration since the price was last changed. Thus, some firms have not changed their prices in a long while. In this situation, the best policy to minimize price dispersion is PT. Of course, with commitment, PT is optimal in the New-Keynesian model. But the analysis does not rely on a PTO in the social loss function.

It is open to question whether incorporating additional long-run rigidities in the product market would yield a PTO in the social loss function in New-Keynesian and hybrid cases. Similarly, rigidities in input or financial markets suggest additional important reasons for a nominal price-level anchor. In all of these markets, a long-term price-level target would assist in planning in nominal quantities and would reduce ex post price-level surprises.[29] In financial contracting, this might yield significant benefits for capital market and economic growth. While these reasons for price-level inertia might suggest including a PTO in the social loss function, they are not self-evident when one considers Woodford's (1999c) analysis.

---

28. Svensson (1999d) derives PT for a New-Classical PC under both commitment and discretion. Dittmar and Gavin (2000) and Vestin (2000) show this for a New-Keynesian PC under discretion. In the remaining polar case, a New-Keynesian PC with commitment, it can be shown that PT is optimal. It follows that with loss function (6), the Bank always engages in PT with a hybrid PC.

29. Christiano and Gust (1999) examine optimal policy in a limited-participation model of short-run financial frictions. Smith (1994) examines PT in an overlapping-generations model with assets. Neither paper derives loss functions.

Still, there is hope for a PTO: the literature has yet to examine how policy endogenously affects contracting and expectations. For example, the Calvo (1983) staggered-price-setting model is used in the New-Keynesian analysis. Yet it is not clear that this model of price setting is optimal in both IT and PT worlds. Similarly, wage and financial contracts may display quite different forms under different policy regimes. It seems reasonable that IT would generate greater inflation inertia, whereas PT would generate greater price-level inertia in aggregate supply. The same reasoning should also apply to how expectations evolve. Coulombe (1998) argues convincingly that a PT regime is much more information-efficient for intertemporal comparisons. Convinced that PT is a credible regime, people would happily switch to the easier price-level-stationary rule of thumb from a price-level-difference (inflationary) stationary rule of thumb. To our minds, the key outstanding issue is gaining a better understanding of how expectations, rigidities, and policy interact to affect welfare.

6 Conclusion: Price-Level Targeting Is a Good Idea When . . .

This paper examines price-level targeting (PT) versus inflation targeting (IT) policy in the standard theoretical framework in which the Bank minimizes a quadratic loss function subject to a log-linear Phillips curve. This linear-regulator framework has recently been put on a firm footing by Rotemberg and Woodford (1997) and Woodford (1999c), who show that it can be derived from a general-equilibrium model with optimizing consumers and monopolistically competitive firms. It is also the framework used for much of applied monetary policy analysis, so it is appropriate for this primer on PT.

6.1 In the standard set-up when . . .

Within the linear-regulator framework, the results for and against PT depend on whether: (i) the Bank can commit to policy or acts under discretion, (ii) expectations in the PC are forward-looking or predetermined, (iii) output persistence in the PC is endogenous or exogenous, (iv) the Bank can target current variables or forecast variables, and (v) the Bank can be assigned alternative loss functions, or not.

The most striking result in the literature is that PT is optimal in the New-Keynesian case with commitment. However, this is a very special case, where expectations are strictly forward-looking. Introducing even a small weight on predetermined expectations in a hybrid PC upsets this result. The general result is that IT is optimal when both predetermined and forward expectations are present in the PC. Inflation targeting is optimal in the sense that optimal commitment policy displays price-level drift. As the weight on the forward-looking component increases, the optimal policy displays less drift. If the Bank cannot target current variables because prices are pre-determined, then even in the New-Keynesian case, PT is not optimal.

The other striking result in the literature is that PT can yield a "free lunch" when the Bank cannot commit to future policy. The free-lunch result indicates that the Bank with discretion should not be assigned the social loss function but instead a loss function that incorporates history. The reason is that the social loss function has an inflation target objective (ITO) and, hence, ignores history under discretion. Hence, the Bank foregoes the ability to condition expectations for price stability. This inability can lead to very poor outcomes. In the New-Classical case, as the endogenous persistence from lagged output increases, the variance of inflation becomes very large and eventually infinite.

When the outcome is bad under discretion, it is better to assign the Bank a PTO. With a PTO, the Bank chooses a PT policy that lowers the variance of inflation while maintaining the variance of output. This free-lunch result obtains because PT acts to condition expectations to prevent output variability from rising. From another perspective, using a PTO is welfare-improving when it yields a PT policy that better mimics the commitment solution. The relative performance of the PTO loss function to the ITO loss function improves the greater is the persistence in the output gap.

However, free lunch may not be only for a PTO loss function, but also for other objective functions that constrain the Bank to consider history. Adding interest rate smoothing or nominal-income objectives into the loss function may improve welfare. But if a range of objectives is feasible, a thorny question arises: why can the Bank commit to alternative objectives but not to the commitment solution? Both commitment and discretion solutions are methodologically problematic. However, both solution concepts point to PT as a good idea when expectations are sufficiently forward-looking.

6.2 In the framework when . . .

The case for PT is very strong if the true social loss function contained a quadratic PTO. Then optimal commitment policy is price-stationary for loss function (6) and other variants that do not also contain a nominal interest rate objective. Such a PTO might well follow from price-level inertia arguments. However, as discussed in the previous section, this has yet to be proven.

The case for weighing the nominal interest rate in the social loss function has been made by Woodford (1999c). When expectations are forward-looking, this objective provides another reason for conditioning expectations for lower inflation by correcting price-level drift. In fact, in the New-Keynesian case, the optimal policy with this objective is to overcorrect shocks. Though not optimal, PT is substantially in the right direction when this objective has importance.

Important arguments for and against PT arise from constraints on the choice variables. Lags in the control of target variables would seem to favour IT, as indicated by Result 4. However, much more work has to be done in this area, since the issues of instruments and controllability have yet to be integrated into the PT literature. An important restriction on policy that has received some analysis in the PT literature is the zero constraint on the nominal interest rate. Such a constraint is of concern because a fixed price-level target is particularly attractive. But this would not leave much room for nominal interest rates to fall, and would leave the economy susceptible to a liquidity trap.

A zero constraint on the nominal interest rate is difficult to model in the linear regulator framework. This constraint introduces an asymmetry into the problem, and we conjecture that optimal commitment policy would have to be asymmetric. In the New-Keynesian case, optimal policy should respond aggressively to counter deflation shocks and overcorrect so that the price level eventually increases. This conditions agents to expect inflation following such a negative shock and not contract output. Such a policy would limit the possibility of liquidity traps.

The possibility of overcorrecting deflationary shocks when expectations are strictly forward-looking suggests that with a hybrid PC an appropriate response would be to restore the fixed price level. Maintaining a fixed price level would be a good clear rule that helps avoid the more serious deflationary shocks.[30] Combined with occasional judicious fiscal policy, a liquidity trap may be avoided.[31] Alternatively, the possibility of a liquidity trap could be greatly reduced by adopting a slowly growing price-level

30. Coulombe (1998) points out that credible PT, in fact, conditions agents to expect future inflation and this results in negative real interest rates even though nominal rates cannot fall below zero. Wolman (1998) finds that PT rules have this advantage over IT rules in simulation analysis of forward-looking models. Reifschneider and Williams (1999) examine the properties of rules in situations where the nominal interest zero bound is likely to bind.

31. Woodford (1999b) argues that the possibility of a liquidity trap is an important issue because of the possibility of self-fulfilling deflationary traps. Then even open market operations and attempts to deflate the exchange rate through purchases of foreign exchange are impotent. Fiscal policy commitments, however, can preclude such undesirable traps.

target of 1 to 2 per cent, as shown by Reifschneider and Williams. Of course, this alternative policy does not have the appeal and simplicity of a fixed price-level target.

Another important asymmetry that should be included in the analysis is an asymmetric PC. We conjecture that it would introduce an asymmetry in policy and upset the New-Keynesian result that PT is optimal. To condition expectations against deflationary shocks, policy should overcorrect such shocks so that the price level eventually increases. Again, this suggests that PT may be a good compromise policy when expectations are only partially forward-looking in a hybrid PC. Maclean and Pioro's (2001) simulations of the Quarterly Projection Model, which contains an asymmetric PC, support PT.

6.3 When more is known about . . .

We have omitted many important issues. Foremost is the extent of downward nominal rigidity, examined by Beaudry and Doyle (2001) and Farès and Lemieux (2001) in this volume. Also in this volume, Maclean and Pioro analyze the important issue of credibility and the transition from IT to PT. Once established, PT may be a good policy, but conditioning expectations for PT may involve a costly transition. Finally, we have omitted open economy issues, as has the rest of the non-simulation literature.

In an open economy, first observe that if two countries are PT, then they are effectively targeting the long-run exchange rate through the purchasing-power-parity condition. Thus, the effect of policy depends on the international regime. Its success or failure hinges on the degree of cooperation across targeting nations, in much the same way that more formal exchange rate arrangements, such as the gold standard, Bretton Woods, or the European Monetary System, have worked. Second, exchange rate shocks are a major source of both real and nominal disturbances. Typically, nominal exchange rates adjust faster than prices; such nominal rigidities can have substantial sector-based shocks. Does PT lead to greater or less short-run variability in the open sectors of the economy? Third, in a world of increasing capital mobility, how does opening the economy affect the possibility of targeting the nominal interest rate? Is it possible for a Bank to get all of them—the interest rate, the exchange rate, and the price level—right?

PT is rarely a perfect policy. The new literature on PT shows that the case for anchoring the price level to a long-run price-level target is a good idea in the standard set-up when expectations are forward-looking and current target variables are controllable. Outside the standard set-up, a number of additional expectations-based arguments could potentially be marshalled in support of PT. In conjunction with the conventional arguments for PT discussed in the introduction, the overall case for PT now seems quite promising compared with the conventional wisdom of only a few years ago. However, as our discussion indicates, we believe much more research is needed before PT can be confidently considered worth implementing.

Appendix

This appendix shows that the commitment solution for the hybrid PC yields an optimal inflation path described by a difference equation (2).

Starting with the optimal inflation-rule equation (10) and the PC, we derive a second-order difference equation in the output gap, x_t.

$$\pi^* - \left(\frac{\lambda}{\kappa}\right)\left[x_t - \phi E_{t-1}x_t\right] + \left(\frac{\lambda}{\kappa}\right)(1-\phi)x_{t-1} = \tag{A1}$$

$$\kappa x_t + \phi\left[\pi^* - \left(\frac{\lambda}{\kappa}\right)(1-\phi)E_{t-1}x_t + \left(\frac{\lambda}{\kappa}\right)(1-\phi)x_{t-1}\right]$$

$$+ (1-\phi)\beta\left[\pi^* - \left(\frac{\lambda}{\kappa}\right)(1-\phi)E_t x_{t+1} + \left(\frac{\lambda}{\kappa}\right)(1-\phi)x_t\right] + u_t$$

Equation (A1) can be written as:

$$x_t = A\pi^* + A_0 E_t x_{t+1} + A_1 x_{t-1} + A_2 E_{t-1}x_t + A_3 u_t, \tag{A2}$$

where

$$A \equiv \frac{(1-\beta)(1-\phi)\kappa}{\Delta};\ A_0 \equiv \frac{\beta\lambda(1-\phi)^2}{\Delta};\ A_1 \equiv \frac{\lambda(1-\phi)^2}{\Delta};$$

$$A_2 \equiv \frac{\lambda\phi(2-\phi)}{\Delta};\ A_3 \equiv \frac{\kappa}{\Delta}$$

and

$$\Delta \equiv \kappa^2 + \lambda[1 + \beta(1-\phi)^2] < 0.$$

Note that $0 < A, A_0, A_1, A_2 < 1$, and $A_3 < 0$.

Equation (A2) can be used to find the forecasts $E_{t-1}x_t$ and $E_t x_{t+1}$, following methods detailed in Blanchard and Fischer (1989, 261–66). These forecasts, along with equation (A2), yield a solution for x_t, in the form

$$x_t = \alpha\pi^* + v x_{t-1} + A_4 \rho u_{t-1} + \gamma\varepsilon_t, \tag{A3}$$

where

$$\alpha \equiv \frac{A(1 + A_2\beta v / A_0)}{(1 - vA_0)(1 - \beta\rho v)} > 0,\ A_4 \equiv A_3\frac{1 + \beta v A_2 / A_0}{(1 - vA_0)(1 - \beta\rho v)} < 0;$$

$$\gamma \equiv \frac{A_3}{(1 - vA_0)(1 - \beta v \rho)} < 0,$$

and v satisfies the characteristic equation $v + \frac{1}{\beta v} = \frac{1 - A_2}{A_0}$.

To verify that v is less than 1 in absolute value, note that the term

$$\frac{1 - A_2}{A_0}$$

is increasing in ϕ. At $\phi = 0$, it equals

$$\frac{\kappa^2}{\lambda\beta} + \frac{(1 + \beta)}{\beta} > \frac{(1 + \beta)}{\beta}.$$

The equation

$$f(z) = z + \frac{1}{\beta z}$$

is decreasing, then increasing over the interval $(0, \infty)$, reaching a minimum at z^*, where

$$z^* = \left(\frac{1}{\beta}\right)^{1/2} \text{ and } 1 < z^* < \frac{1}{\beta}.$$

Since

$$f(1) = f(1/\beta) = \frac{1 + \beta}{\beta},$$

it follows that one root must be greater than $\frac{1}{\beta}$, one less than 1, in order to satisfy the condition

$$f(z) = 1 - \frac{A_2}{A_0}.$$

The root that is greater than 1 is

$$\frac{1}{\beta v}.$$

Further, as $\frac{1 - A_2}{A_0} > 0$, the root less than 1 is greater than zero, $0 < v < 1$.

Using the decision rule,

$$\pi_t = \pi^* - \frac{\lambda}{\kappa}(x_t - \phi E_{t-1}x_t) + \frac{\lambda}{\kappa}(1-\phi)x_{t-1}$$

and equation (A3), we can arrive at a solution for the inflation rate in terms of past inflation, the shocks, and a constant. We have:

$$\pi_t = \pi^* - \frac{\lambda}{\kappa}\Big[\alpha\pi^* + vx_{t-1} + A_4\rho u_{t-1} + \gamma\varepsilon_t$$

$$- \phi(\alpha\pi^* + vx_{t-1} + A_4\rho u_{t-1})\Big] + \frac{\lambda}{\kappa}(1-\phi)x_{t-1}.$$

Collecting terms,

$$\pi_t = \left[1 - \frac{\lambda}{\kappa}(1-\phi)\alpha\right]\pi^* + \frac{\lambda}{\kappa}(1-\phi)(1-v)x_{t-1} \tag{A4}$$

$$-\frac{\lambda}{\kappa}(1-\phi)A_4\rho u_{t-1} - \frac{\lambda}{\kappa}\gamma\varepsilon_t.$$

Lagging equation (A4) one period, we solve for x_{t-2},

$$x_{t-2} = \frac{-\pi^*\left[1 - \frac{\lambda}{\kappa}(1-\phi)\alpha\right] + \pi_{t-1} + \frac{\lambda}{\kappa}(1-\phi)A_4\rho u_{t-2} + \frac{\lambda}{\kappa}\gamma\varepsilon_{t-1}}{\frac{\lambda}{\kappa}(1-\phi)(1-v)}.$$

We then use equation (A3), lagged one period, and this solution for x_{t-2} to write x_{t-1} in terms of π_{t-1}, u_{t-2}, ε_{t-1}, and a constant:

$$x_{t-1} = \frac{\pi^*\left[-v + \frac{\lambda}{\kappa}(1-\phi)\alpha\right] + v\pi_{t-1} + \frac{\lambda}{\kappa}(1-\phi)A_4\rho u_{t-2} + [(1-\phi) + v\phi]\frac{\lambda}{\kappa}\gamma\varepsilon_{t-1}}{\frac{\lambda}{\kappa}(1-\phi)(1-v)}.$$

Using this solution in equation (A4), we can arrive at a solution for the inflation rate:

$$\pi_t = (1-v)\pi^* + v\pi_{t-1} - (1-\phi)A_4\frac{\lambda}{\kappa}\rho(u_{t-1} - u_{t-2}) \tag{A5}$$

$$-\frac{\gamma\lambda}{\kappa}\varepsilon_t + \frac{\gamma\lambda}{\kappa}[(1-\phi) + v\phi]\varepsilon_{t-1}$$

Inspecting (A5), we see that the inflation process is stationary as $0 < v < 1$. Equation (A5) takes the form of equation (2):

$$\pi_t = a\pi_{t-1} + (1-a)\pi^* + b\rho(u_{t-1} - u_{t-2}) + c\varepsilon_t - d\varepsilon_{t-1},$$

where $a = v$,

$$b = -(1-\phi)A_4\frac{\lambda}{\kappa} > 0,\ c = \frac{\gamma\lambda}{\kappa} > 0, \text{ and}$$

$$d = -\frac{\gamma\lambda}{\kappa}[(1-\phi) + a\phi] > 0.$$

The price-level process is not stationary, except in the New-Keynesian case, $\phi = 0$, where $b = c = d$ and the equation can be expressed like (1').

References

Backus, D.K. and J. Driffill. 1986. "The Consistency of Optimal Policy in Stochastic Rational Expectations Models." Centre for Economic Policy Research Discussion Paper No. 124.

Beaudry, P. and M. Doyle. 2001. "What Happened to the Phillips Curve in the 1990s in Canada?" In *Price Stability and the Long-Run Target for Monetary Policy,* 51–82. Proceedings of a seminar held by the Bank of Canada, June 2000. Ottawa: Bank of Canada.

Berg, C. and L. Jonung. 1999. "Pioneering Price Level Targeting: The Swedish Experience 1931–1937." *Journal of Monetary Economics* 43 (3): 525–51.

Bernanke, B.S., T. Laubach, F.S. Mishkin, and A. Posen. 1999. *Inflation Targeting: Lessons from the International Experience*. Princeton: Princeton University Press.

Black, R., D. Coletti, and S. Monnier. 1998. "On the Costs and Benefits of Price Stability." In *Price Stability, Inflation Targets, and Monetary Policy,* 303–42. Proceedings of a conference held by the Bank of Canada, May 1997. Ottawa: Bank of Canada.

Black, R., T. Macklem, and D. Rose. 1998. "On Policy Rules for Price Stability." In *Price Stability, Inflation Targets, and Monetary Policy,* 411–61. Proceedings of a conference held by the Bank of Canada, May 1997. Ottawa: Bank of Canada.

Blanchard, O.J. and S. Fischer. 1989. *Lectures on Macroeconomics*. Cambridge, MA: MIT Press.

Calvo, G.A. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12 (3): 383–98.

Chari, V.V., P.J. Kehoe, and E.C. Prescott. 1989. "Time Consistency and Policy." In *Modern Business Cycle Theory,* edited by R.J. Barro, 265–305. Cambridge, MA: Harvard University Press.

Christiano, L.J. and C.J. Gust. 1999. "Taylor Rules in a Limited Participation Model." NBER Working Paper No. 7017.

Clarida, R., J. Galí, and M. Gertler. 1999. "The Science of Monetary Policy: A New Keynesian Perspective." *Journal of Economic Literature* 37 (4): 1661–707.

Coulombe, S. 1998. "The Intertemporal Nature of Information Conveyed by the Price System." In *Price Stability, Inflation Targets, and Monetary Policy,* 3–28. Proceedings of a conference held by the Bank of Canada, May 1997. Ottawa: Bank of Canada.

Currie, D. and P. Levine. 1993. *Rules, Reputation and Macroeconomic Policy Coordination*. Cambridge: Cambridge University Press.

Dehejia, V.H. and N. Rowe. 2000. "Macroeconomic Stabilisation: Fixed Exchange Rates vs Inflation Targeting vs Price Level Targeting." Carleton University. Unpublished.

Dittmar, R., W.T. Gavin, and F.E. Kydland. 1999. "The Inflation-Output Variability Tradeoff and Price-Level Targets." *Federal Reserve Bank of St. Louis Review* 81 (1): 23–31.

Dittmar, R. and W.T. Gavin. 2000. "What Do New-Keynesian Phillips Curves Imply for Price-Level Targeting?" *Federal Reserve Bank of St. Louis Review* 82 (2): 21–30.

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

Farès, J. and T. Lemieux. 2001. "Downward Nominal-Wage Rigidity: A Critical Assessment and Some New Evidence for Canada." In *Price Stability and the Long-Run Target for Monetary Policy,* 3–31. Proceedings of a seminar held by the Bank of Canada, June 2000. Ottawa: Bank of Canada.

Feldstein, M. 1997. "Costs and Benefits of Going from Low Inflation to Price Stability." In *Reducing Inflation: Motivation and Strategy,* edited by C. Romer and D. Romer, 123–56. Chicago: Chicago University Press.

Fillion, J.-F. and R. Tetlow. 1994. "Zero-Inflation or Price-Level Targeting? Some Answers from Stochastic Simulations on a Small Open-Economy Macro Model." *Economic Behaviour and Policy Choice under Price Stability,* 129–66. Ottawa: Bank of Canada.

Fischer, S. 1994. "Modern Central Banking." In F. Capie, C. Goodhart, S. Fischer, and N. Schnadt, *The Future of Central Banking: The Tercentenary Symposium of the Bank of England*, 262–308. Cambridge: Cambridge University Press.

Fuhrer, J. and G. Moore. 1995. "Inflation Persistence." *Quarterly Journal of Economics* 110 (1): 127–59.

Haldane, A.G. and C.K. Salmon. 1995. "Three Issues on Inflation Targets." In *Targeting Inflation*, edited by A. Haldane, 170–201. A conference of central banks on the use of inflation targets organised by the Bank of England, March 1995. London: Bank of England.

Hall, R.E. 1984. "Monetary Strategy with an Elastic Price Standard." In *Price Stability and Public Policy*, 137–59. Kansas City: Federal Reserve Bank of Kansas City.

Hall, R.E. and N.G. Mankiw. 1994. "Nominal Income Targeting." In *Monetary Policy*, edited by N.G. Mankiw, 29: 79–93. NBER Studies in Business Cycles. Chicago: University of Chicago Press.

Howitt, P. 2000. "Learning About Monetary Theory and Policy." Brown University. Photocopy.

Jensen, H. 1999. "Targeting Nominal Income Growth or Inflation?" Centre for Economic Policy Research Discussion Paper No. 2341.

Kiley, M.T. 1998. "Monetary Policy under Neoclassical and New-Keynesian Phillips Curves, with an Application to Price Level and Inflation Targeting." Finance and Economics Discussion Series, Discussion Paper No. 1998–27. Federal Reserve Board.

King, R.G. and A.L. Wolman. 1999. "What Should the Monetary Authority Do When Prices Are Sticky?" In *Monetary Policy Rules*, edited by J.B. Taylor, 349–98. Chicago: University of Chicago Press.

Konieczny, J.D. 1994. "The Optimal Rate of Inflation: Competing Theories and Their Relevance to Canada." In *Economic Behaviour and Policy Choice Under Price Stability*, 1–35. Proceedings of a conference held by the Bank of Canada, October 1993. Ottawa: Bank of Canada.

Kydland, F.E. and E.C. Prescott. 1977. "Rules Rather than Discretion: The Inconsistency of Optimal Plans." *Journal of Political Economy* 85 (3): 473–91.

Lebow, D.E., J.M. Roberts, and D.J. Stockton. 1992. "Economic Performance under Price Stability." Working Paper No. 125. Division of Research and Statistics, Federal Reserve System.

Lucas, R.E., Jr. 1973. "Some International Evidence on Output-Inflation Tradeoffs." *American Economic Review* 63 (3): 326–34.

Maclean, D. and H. Pioro. 2001. "Price-Level Targeting—The Role of Credibility." In *Price Stability and the Long-Run Target for Monetary Policy*, 153–85. Proceedings of a seminar held by the Bank of Canada, June 2000. Ottawa: Bank of Canada.

McCallum, B.T. 1994. "A Semi-Classical Model of Price-Level Adjustment." *Carnegie-Rochester Conference Series on Public Policy* (41): 251–84.

———. 1995. "Two Fallacies Concerning Central-Bank Independence." *American Economic Review* 85 (2): 207–11.

———. 1996. "Crucial Issues Concerning Central Bank Independence." NBER Working Paper No. 5597.

———. 1999. "Issues in the Design of Monetary Policy Rules." In *Handbook of Macroeconomics*, edited by J.B. Taylor and M. Woodford (1C) Part 7, Chapter 23. Amsterdam: Elsevier Science.

Mishkin, F.S. 2001. "Issues in Inflation Targeting." In *Price Stability and the Long-Run Target for Monetary Policy*, 203–22. Proceedings of a seminar held by the Bank of Canada, June 2000. Ottawa: Bank of Canada.

Parkin, M. 2001. "What Have We Learned About Price Stability?" In *Price Stability and the Long-Run Target for Monetary Policy*, 223–59. Proceedings of a seminar held by the Bank of Canada, June 2000. Ottawa: Bank of Canada.

Reifschneider, D. and J.C. Williams. 1999. "Three Lessons for Monetary Policy in a Low Inflation Era." Finance and Economics Discussion Series 1999–44. Federal Reserve Board.

Roberts, J.M. 1995. “New Keynesian Economics and the Phillips Curve.” *Journal of Money, Credit and Banking* 27 (4) Part 1: 975–84.

Rogoff, K. 1985. “The Optimal Degree of Commitment to an Intermediate Monetary Target.” *Quarterly Journal of Economics* 100 (4): 1169–89.

Rotemberg, J.J. and M. Woodford. 1997. “An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy.” *NBER Macroeconomics Annual* (1997): 297–346.

———. 1999. “Interest Rate Rules in an Estimated Sticky Price Model.” In *Monetary Policy Rules*, edited by J.B. Taylor, 57–119. Chicago: University of Chicago Press.

Smith, B.D. 1994. “Efficiency and Determinacy of Equilibrium under Inflation Targeting.” *Economic Theory* 4 (3): 327–44.

Svensson, L.E.O. 1997a. “Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets.” *European Economic Review* 41 (6): 1111–46.

———. 1997b. “Optimal Inflation Targets, ‘Conservative’ Central Banks, and Linear Inflation Contracts.” *American Economic Review* 87 (1): 98–114.

———. 1997c. “Price-Level Targeting vs. Inflation Targeting: A Free Lunch?” Working Paper No. 614. Stockholm University.

———. 1999a. “Inflation Targeting as a Monetary Policy Rule.” *Journal of Monetary Economics* 43 (3): 607–54.

———. 1999b. “Inflation Targeting: Some Extensions.” *Scandinavian Journal of Economics* 101 (3): 337–61.

———. 1999c. “How Should Monetary Policy Be Conducted in an Era of Price Stability?” In *New Challenges for Monetary Policy*, 195–259. A symposium sponsored by the Federal Reserve Bank of Kansas City, August 1999. Federal Reserve Bank of Kansas City.

———. 1999d. “Price-Level Targeting versus Inflation Targeting: A Free Lunch?” *Journal of Money, Credit and Banking* 31 (3) Part 1: 277–95.

Svensson, L.E.O. and M. Woodford. 1999. “Implementing Optimal Policy Through Inflation-Forecast Targeting.” November 1999. Princeton University. Unpublished. < URL: http://www.princeton.edu/~woodford/ > (February 2000).

Vestin, D. 2000. “Price-Level Targeting versus Inflation Targeting in a Forward-Looking Model.” Working Paper Series ISSN1402–1903, No. 106, Sveriges Riksbank, Stockholm.

Williams, J.C. 1999. “Simple Rules for Monetary Policy.” Finance and Economics Discussion Series Paper No. 1999–12. Federal Reserve System.

Wolman, A.L. 1998. “Staggered Price Setting and the Zero Bound on Nominal Interest Rates.” *Economic Quarterly* 84 (4): 1–24.

Woodford, M. 1999a. “Commentary: How Should Monetary Policy Be Conducted in an Era of Price Stability?” In *New Challenges for Monetary Policy*, 277–316. Kansas City: Federal Reserve Bank of Kansas City. < URL: http://www.princeton.edu/~woodford/ >.

———. 1999b. “Price-Level Determination under Interest-Rate Rules.” April 1999. Princeton University. Unpublished. < URL: http://www.princeton.edu/~woodford/ >.

———. 1999c. “Inflation Stabilization and Welfare.” June 1999. Princeton University. Unpublished. < URL: http://www.princeton.edu/~woodford/ >.

———. 1999d. “Optimal Monetary Policy Inertia.” NBER Working Paper No. 7261.

———. 2000. “Pitfalls of Forward-Looking Monetary Policy.” *American Economic Review* 90 (2): 100–4.

Discussion

Jean Boivin

Introduction

In this interesting paper, Barnett and Engineer investigate the conditions under which price-level targeting (PT) might be desirable. The focus of their analysis is on the “new” literature supporting PT on the basis that it provides a useful anchor for expectations.

Traditionally, the debate between inflation targeting (IT) and PT revolved around an apparent trade-off between the short-run and long-run variability of inflation and output.[1] Recently, however, a new argument (which I will refer to as the “new argument”) in favour of PT arose from a better understanding of the implications of a forward-looking private sector in an intertemporal model. Essentially, the idea is that by favourably conditioning inflation expectations, PT could lead to an improved trade-off today between inflation and output, thus resulting in a better equilibrium with lower inflation variability, as well as potentially lower output variability. This has implications in terms of defining the optimal commitment solution, as well as potentially improving on the discretionary solution (the “free-lunch” result).

This paper’s main contribution is an assessment of the robustness of this argument to the modelling ingredients. In particular, the literature survey shows how assumptions about the private sector expectations entering the Phillips curve (PC) impinge on the desirability of PT. While their analysis draws on these existing results, the authors also derive useful new results

1. As discussed in Barnett and Engineer’s paper, proponents of PT emphasized that an IT rule would lead to unbounded long-run price-level variability (since the implied price-level path is I(1)), whereas proponents of IT emphasized the higher short-run variability of output that would result from targeting the price level.

from a "hybrid" PC that blends the different modelling features considered in the literature in a unified framework. The key result from their analysis is that under commitment and for a "standard" social loss function, a stationary price level is desirable only for a *purely* forward-looking (New-Keynesian) specification of the PC. Otherwise, the optimal solution involves a drifting price level (I(1)).

The focus of my discussion is first on the interpretation of this result, especially in terms of its policy implications, and second, on the importance of the assumed social loss function in driving the results. As in their paper, attention is essentially given to the commitment case. But first, it is useful to review the intuition behind both the new argument for PT and the key result from the paper.

Contradictory results have been reached in the literature concerning the properties of the optimal price-level solution under commitment. For instance, Svensson (1999) found that the optimal commitment solution for a model based on a New-Classical PC involved a drifting price level, as would be implemented by an IT rule.[2] On the other hand, Woodford (1999b, 2000) found that for a New-Keynesian specification of the PC, the optimal commitment solution involved a stationary price process, as would be implemented through PT. Why do different specifications of the PC lead to drastically different solutions for the price level?

The role of private sector expectations in determining the nature of the price path under the optimal equilibrium can be easily seen by considering the properties of two standard specifications of the PC. The two polar cases of the hybrid PC considered by Barnett and Engineer are[3]:

$$\pi_t = \kappa x_t + \beta E_{t-1}\pi_t + u_t \tag{1}$$

$$\pi_t = \kappa x_t + \beta E_t\pi_{t+1} + u_t \tag{2}$$

The first specification corresponds to the Lucas supply curve and is used in the context of PT by Svensson (1999) and Dittmar et al. (1999). The second specification corresponds to the New-Keynesian PC and can be derived, as in Calvo (1983), from a model of optimal staggered-price setting. It is used in the present context, among others, by Woodford (1999b, 2000), Dittmar and Gavin (2000), and Vestin (2000). As Barnett and Engineer emphasize,

2. According to Barnett and Engineer's definition, an IT rule leads to a long-run price-level drift. This issue will be discussed further.

3. I am using the same notation as in their paper. Note that I am considering the simplest case; in particular, there is no persistence. That is sufficient, however, to emphasize the intuition in the commitment case.

although deceptively similar, these two specifications have quite different implications for the desirability of a PT rule under commitment.

Consider the effect of a cost-push shock hitting these two specifications. In the “Lucas supply” curve, i.e., the New-Classical case, a combination of increased inflation and decreased output must totally offset this shock, since $E_{t-1}\pi_t$ is predetermined. Because there is no policy response that can affect the current trade-off between output and inflation, the optimal response is to accommodate the inflationary shock. The inflation cost is fully absorbed in the current period and cannot be reduced.[4] Since society does not care about price-level deviations per se, the optimal solution involves a drifting price level.

In the New-Keynesian PC, $E_t\pi_{t+1}$ is based on information about the current cost-push shock, as well as the *current expected policy response*. If, in the next period, the central bank can credibly commit to a policy of lower inflation, it can favourably alter the current sacrifice ratio; by promising to offset the inflationary shock next period, part of u_t in equation (2) is absorbed by a decrease in expected inflation. For this to be an equilibrium, the promise has to be fulfilled. The optimal commitment policy for offsetting the inflationary shock thus involves trading today’s higher inflation at a lower output cost, in exchange for future lower inflation, which would lead to a stationary price level.

As just illustrated, depending on the nature of the private sector expectations, opposite results are reached concerning the desirability of a (trend) stationary price level. The relevant question then becomes: which characterization of the private sector expectations is more relevant in practice? One way to get at this is to determine what happens when the PC has both forward-looking and predetermined expectations. Significantly, Barnett and Engineer were able to derive the optimal commitment solution in the case of a hybrid PC that includes both types of expectations. No micro-foundations for this kind of PC are provided, but as I will discuss, Woodford (1999a) provides an example of pricing decisions timing that would yield a similar PC.

Barnett and Engineer conclude from their Result 2 that, assuming a quadratic loss function in output and inflation, the optimal solution under commitment involves a stationary price *only* in the polar case of purely forward-looking agents. As soon as there is a small fraction of predetermined expectations, the optimal solution calls for some price-level drifts.

---

4. Obviously, as suggested by Barnett and Engineer’s Result 1, if the loss function had a zero weight on output-gap deviation, there would be no conflict between objectives: inflation would always be equal to its target and so would the price level.

The intuition behind this result is fairly straightforward. Let us first consider the New-Keynesian PC. In this case, undoing inflationary shocks helps condition expectations of future inflation, which leads to improved welfare through lower variability of inflation and output. If agents with predetermined expectations are introduced, *and since price-level deviations do not directly enter the loss function*, the central bank will face a trade-off. Completely undoing the inflationary shocks while still benefiting the forward-looking agents is costly for those who do not take the expected policy into account in their pricing decisions. Faced with this trade-off, the central bank finds it optimal to let part of the shocks have a permanent effect on the price level. Ultimately, if the fraction of a backward-looking agent goes to 1, the benefits from partially undoing the shocks will go to 0, which gives us the result obtained in the other polar case of the New-Classical PC.

Based on their Result 2, it thus appears that a stationary price-level process will not generally be part of the optimal solution.

Result 2 is really a characterization of the price-level properties under the optimal commitment equilibrium. An important question is how we should translate Result 2 in terms of a policy prescription. Or, what can we say about the question asked in the paper's title: when is price-level targeting a good idea?

The authors' answer seems to be that since a stationary price path is optimal only in the "non-generic" case of the New-Keynesian PC, PT is not a good idea under the assumptions of Result 2. In fact, any other specifications of the PC would lead to some price-level drift, suggesting that IT should be favoured over PT.

It is important to note, however, that Barnett and Engineer use a very general definition of IT. In fact, they define it as "a policy that systematically responds to deviations of the inflation rate from the target inflation rate in a way that effectively yields long-run price-level drift" (see page 105). While this clearly includes *pure inflation targeting*, i.e., when there is no attempt to partially undo past inflationary shocks, it also includes "hybrid" policies that target the inflation rate as well as, partially, the price level. So IT includes a continuum of policy reaction functions, involving (more or less) an undoing of the shocks, and as a result (more or less), price-level drifts. Consequently, IT is a much broader class than PT, and thus it might not be surprising that it is the generic optimal solution.

Since PT can be seen as a limiting case of the IT class, I would argue that a relevant and "fairer" comparison should be between PT and *one particular element* of the IT class. Moreover, it is not clear that all of the policy rules in the IT class are relevant in practice. In fact, for credibility, accountability, and transparency issues, the feasible choices of rule might be between only the simple PT and *pure* IT. At least one could argue that they are the most commonly discussed forms of policies (for instance, IT in practice as examined in Bernanke et al. 1999).

If these constitute the only two feasible policy alternatives, then PT would not be optimal solely for the New-Keynesian PC. In fact, for a model having a small fraction of agents with predetermined expectations, PT should be closer to the optimal, non-feasible policy response.

Based on this discussion, it is not clear that we can draw practical policy prescriptions directly from Result 2. Clearly, if the optimal policy rule is always feasible, PT will be desirable only in a very specific, perhaps even unlikely, case. But the same is true for pure inflation targeting or, for that matter, for any element of the IT class taken in isolation. The important point, however, is that if we see some specific rules as an approximation of an infeasible, optimal general rule, a specific PT rule might be optimal for a much broader class of PC specifications than that suggested by Result 2.

Thus far, I have discussed only the case for PT when the social loss function does not include a price-level objective (what the authors call a PTO). If PT were found to be optimal in this kind of framework, the argument in favour of it would obviously be strong.

However, this specification of the loss function is not theoretically consistent with all the specifications of the PC encompassed by the hybrid PC considered in this paper. As Barnett and Engineer argue, their benchmark loss function can still be justified on the grounds that it has been widely used, albeit on an ad hoc basis, in this literature. But once we recognize that the loss function might be missing social costs that are embodied in the model's assumptions, it is not clear that the relevant optimal commitment solution derived from the hybrid PC should generally involve a drifting price level.

More precisely, unlike the assumptions used in the derivation of Result 2, the social loss function should depend on the specification of the PC. For instance, Woodford's (1999a) derivation of the quadratic loss function relies on the New-Keynesian specification of the PC, or more precisely, it assumes a staggered price setting à la Calvo. On the other hand, Woodford (1999a) obtains a different social loss function—which depends on unanticipated inflation as of time $t-1$ as well as actual inflation—when changing the timing of price changes. In fact, the particular change in the timing that he considers yields a PC very similar to the hybrid PC considered by Barnett and Engineer.[5] Under Woodford's hybrid PC and

5. The difference is essentially that, apart from the persistence in the output gap, the term $\kappa(x_t - \delta x_{t-1})$ in equation (7) of the paper should be multiplied by $(1-\phi)$.

implied loss function, *price-level stability is generally optimal*.[6] Barnett and Engineer acknowledge this point, but it is important to realize that it does upset their Result 2, which, given their hybrid PC, is derived from an ad hoc specification of the loss function.

The question of the "right" specification of the loss function really takes us back to the original debate mentioned in the introduction of this discussion. Ultimately, the desirability of PT relies on the benefits of a stationary price level—and its beneficial effects on private-sector expectations—versus the costs of short-run variability of inflation and unemployment, as well as other costs such as those associated with the higher frequency of deflation periods. Unless these costs and benefits are explicitly included in the model, it appears difficult to conclude whether or not PT is generally desirable.

References

Bernanke, B.S., T. Laubach, F.S. Mishkin, and A. Posen. 1999. *Inflation Targeting: Lessons from the International Experience.* Princeton: Princeton University Press.

Calvo, G.A. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* (12) 3: 383–98.

Dittmar, R., W.T. Gavin, and F.E. Kydland. 1999. "The Inflation-Output Variability Tradeoff and Price-Level Targets." *Federal Reserve Bank of St. Louis Review* 81 (1): 23–31.

Dittmar, R. and W.T. Gavin. 2000. "What Do New-Keynesian Phillips Curves Imply for Price-Level Targeting?" *Federal Reserve Bank of St. Louis Review* 82 (2): 21–30.

Giannoni, M. 2000. "Optimal Interest-Rate Rules in a Forward-Looking Model, and Inflation Stabilization vs Price-Level Stabilization." Working Paper. Federal Reserve Bank. October.

Svensson, L.E.O. 1999. "Price-Level Targeting versus Inflation Targeting: A Free Lunch?" *Journal of Money, Credit and Banking* 31 (3) Part 1: 277–95.

Vestin, D. 2000. "Price-Level Targeting versus Inflation Targeting in a Forward-Looking Model." Working Paper Series ISSN1402–1903, No. 106, Sveriges Riksbank, Stockholm.

Woodford, M. 1999a. "Inflation Stabilization and Welfare." Princeton University. Unpublished. < http://www.princeton.edu/~woodford >.

———. 1999b. "Commentary: How Should Monetary Policy Be Conducted in an Era of Price Stability?" In *New Challenges for Monetary Policy,* 277–316. Federal Reserve Bank of Kansas City. < http://www.princeton.edu/~woodford >.

———. 2000. "Pitfalls of Forward-Looking Monetary Policy." *American Economic Review* 90 (2): 100–4.

6. See also Giannoni (2000) for a comparison of PT and IT when the social loss function is derived from first principle and consistent with the form of the PC.

General Discussion*

Serge Coulombe stated that it is difficult to compare price-level-targeting and inflation-targeting monetary policy regimes in a model without interest rates. He suggested that much of the stabilization effect of price-level targeting would arise from the interaction between movements in both the price level and nominal interest rates. Coulombe also noted that in the model used to assess the price-level-targeting regime, price expectations were a mixture of backward- and forward-looking elements; in his view, in a world of price-level targeting where the price level was trend-stationary and monetary policy was credible, changes in expectations regarding the price level would be instantaneous.

Nicholas Rowe suggested that price-level targeting was optimal even in a neoclassical model if there were both a one-period lag for the effect of monetary policy on aggregate output and a two-period lag for the effect of aggregate demand on wages. In his view, this would be a relatively realistic model.

Jeffrey Fuhrer felt that there was a lack of integration between the loss function of central banks and the behaviour of private agents in the types of models analyzed in the Barnett-Engineer paper. In these models, it is unclear why private agents would dislike inflation, inflation variability, or the variability of prices. Such integration is absolutely necessary, in his view, before one can confidently undertake an analysis of the merits of inflation targeting versus price-level targeting.

* Prepared by Gerald Stuber.

In response to the first comment by Coulombe, Engineer and Barnett noted that in a model featuring both price-level targeting and interest rates in the loss function of the central bank, there would be overshooting of the price level, providing another reason to support price-level targeting. In reference to Fuhrer's comment, it was noted that some recent papers have derived the central bank's loss function to be consistent with the overall characteristics of the model economy.