BANK OF CANADA
BANQUE DU CANADA

Discussion Paper/Document d'analyse
2007-11

Price-Level Targeting and Stabilization Policy: A Review

by Steve Ambler

www.bank-banque-canada.ca

Bank of Canada Discussion Paper 2007-11

October 2007

Price-Level Targeting and Stabilization Policy: A Review

by

Steve Ambler

CIRPÉE, Université du Québec à Montréal

ISSN 1914-0568

Acknowledgements

This paper was written while I was Special Adviser at the Bank of Canada. I would like to thank Robert Amano, Agathe Côté, Serge Coulombe, Tiff Macklem, Stephen Murchison, and John Murray for helpful discussions. All errors are my own responsibility.

Abstract

The author surveys recent articles on the costs and benefits of price-level targeting versus inflation targeting, focusing on the benefits and costs of price-level targeting as a tool for stabilization policy. He reviews papers that examine how price-level targeting affects the short-run trade-off between output and inflation variability by influencing expectations of future inflation. The author looks at the implications of this argument for assigning an objective based on price-level targeting to a central bank that is unable to commit to its future policies. He discusses some recent papers that examine how price-level targeting can help to avoid the zero-bound problem, and papers that examine the incentives created by price-level targeting to change the degree of indexation of private contracts.

JEL classification: E31, E32, E52
Bank classification: Monetary policy framework

Résumé

L'auteur passe en revue les récents articles consacrés à la comparaison des avantages et des coûts respectifs des cibles de niveau des prix et des cibles d'inflation. Il s'attache plus particulièrement à examiner l'utilité des cibles fondées sur le niveau des prix pour une politique de stabilisation. Il analyse des travaux qui tentent de déterminer dans quelle mesure les cibles de niveau des prix jouent un rôle dans l'arbitrage à court terme entre la variabilité de l'inflation et celle de la production en influençant les attentes relatives à l'inflation future. L'auteur s'intéresse aux implications de cette démarche dans l'hypothèse où une banque centrale qui ne peut prendre d'engagement à l'égard de ses orientations à venir choisit un objectif axé sur une cible de niveau des prix. Il évalue aussi des études récentes qui examinent comment les cibles de niveau des prix permettraient d'éviter que les taux d'intérêt atteignent leur limite inférieure de zéro, et d'autres traitant des incitations que créent les cibles de niveau des prix pour modifier le degré d'indexation des contrats privés.

Classification JEL : E31, E32, E52
Classification de la Banque : Cadre de la politique monétaire

1 Introduction

In November 2006 the Bank of Canada and the Government of Canada announced the renewal of the Bank's inflation-control target for a period of five years to the end of 2011. The agreement stipulated that the Bank would continue to aim at keeping CPI inflation at 2 per cent, with a 1 to 3 per cent control range around the target. In its background document on the renewal of the target (Bank of Canada 2006), the Bank announced its intention to lead a research program to address whether and how the monetary policy framework in Canada might be improved. The background document raised two broad sets of questions. The first related to the possibility of lowering the inflation target below 2 per cent. The second related to the potential costs and benefits of replacing the inflation-targeting (IT) regime with a price-level targeting (PLT) regime. An IT regime is defined as a regime in which the central bank aims to keep some measure of inflation, such as CPI inflation, close to a target rate. Under a PLT regime, the central bank's aim is to stabilize the price level around a known target path.[1]

This paper is concerned with the second set of questions. There is a substantial body of research that examines the costs and benefits of PLT compared with those of IT. This paper undertakes a survey of the modern literature on PLT. Traditional analyses[2] focus on the long-run predictability of prices as the main benefit of PLT, while the potential for increased short-run variability of inflation and output are seen as the main drawback. Several recent papers challenge this view, and find that PLT can favourably affect the short-run trade-off between output and inflation variability by affecting expectations of future inflation. The seminal article by Svensson (1999), discussed in detail herein, demonstrates the possibility of a "free lunch" (reduced inflation variability without an increase in output variability) by assigning a target to the central bank with the price level as one of its objectives. Much of the recent literature on PLT centres on analyzing the robustness of this free lunch result in more complex economic environments than the one set out by Svensson. This paper surveys this literature, and aims to assess which conclusions are robust and which questions remain open.

The paper is structured as follows. Section 2 briefly reviews the traditional arguments for and against PLT. Section 3 takes up the "free lunch argument" of Svensson (1999). Section

1. A PLT regime does not necessarily mean that the long-run price level is constant, since the target path may have a positive slope (which determines the long-run rate of inflation). It does mean that the central bank acts to offset deviations of the price level from the target path.
2. The older literature on PLT extends back to Keynes, Fisher, Wicksell, and others. See Duguay (1994) for a cogent summary.

4 examines extensions to the economic environment of Svensson's paper. It discusses the result that in a standard New Keynesian model, with forward-looking price-setters and "cost-push" shocks that introduce a genuine trade-off between inflation and output variability, optimal monetary policy under commitment implies a stationary price level. The section reviews papers that show that some degree of price-level drift (non-stationarity) is optimal in New Keynesian models with price-setters that set their prices using rules of thumb. It discusses Vestin's (2006) result that assigning a price-level target to a central bank that cannot precommit to its future policies can allow it to achieve results that are as good as under commitment. Finally, section 4 reviews papers that examine targeting a moving average of past inflation rates as a compromise between pure PLT and pure IT. Section 5 examines other issues related to PLT. It looks at the effects of a change in monetary regime on the incentive to index contracts, and how that can influence the advantages and costs of PLT; the issue of how to deal with prolonged movements in relative prices; and issues related to the transition between an IT regime and a PLT regime. Section 6 attempts to draw some general conclusions and suggests possible avenues for future research.

2 Traditional Arguments For and Against Price-Level Targeting

The Bank of Canada's current target rate of inflation is 2 per cent. If the annualized rate of inflation is unexpectedly above the 2 per cent target during the current period, then under the Bank's IT regime the target remains at 2 per cent going forward. Under a PLT regime, the inflation target would be reduced to below 2 per cent until the price level itself returned to its original targeted growth path. The effect of the inflation surprise on the price level would be completely offset. Under IT, there is no such offset: a temporary inflation shock leads to a permanent shift in the time path of the price level, and shocks to inflation have a cumulative impact on the price level. As the forecast horizon increases, the forecast-error variance for the price level increases. In the limit, as the forecast horizon goes to infinity, the conditional forecast-error variance for the price level becomes unbounded. The price level is unpredictable at sufficiently long horizons. Under PLT, the conditional forecast-error variance of the price level remains finite at all time horizons, and does not grow monotonically with the length of the horizon.[3]

3. In the technical language of time-series analysis, the price level is "trend stationary" or I(0) under PLT, whereas, under IT, the price level will be permanently affected by a temporary shock: it will be "difference stationary" or I(1).

The long-run predictability of the price level under PLT is precisely the source of the intuitive appeal of this type of monetary regime. It means that the real value of future payments specified contractually in nominal terms is more predictable than under an IT regime, since the conditional variance of the price level increases with the forecast horizon under IT.[4] Under a PLT regime, current prices convey intertemporal information, as argued by Coulombe (1998a,b).

This begs the question as to why individuals sign long-term contracts that stipulate the value of future payments in nominal terms. There is not a strong consensus among economists as to why this is the case, but the prevalence of contracts with fixed nominal payments is not in doubt. Fischer (1994) argues that the benefits of reduced uncertainty concerning the real value of payments could not be very high, given that individuals in the private sector could easily use other means such as indexed bonds and contingent contracts to mitigate the uncertainty without any change in the monetary policy regime. Others infer, on the basis of the same evidence, that the use of these measures by individuals must be economically costly. For example, Howitt (2001) judges that "long-term price-level uncertainty is one of the most serious consequences of inflation, because of its ruinous effects on long-term contracting."[5]

If reduced price-level uncertainty is the main argument that is traditionally invoked in favour of PLT, the traditional argument against PLT is that it must raise the short-run variability of both inflation and output. The logic of this argument is relatively straightforward. In response to a temporary, unexpected increase (decrease) in inflation in a PLT regime, inflation would have to be reduced below (above) its long-run target rate in the short run, in order to move the price level back to its target growth path. The conditional variability of inflation and the price level would be higher than under an IT regime, since, under the latter, monetary policy would merely aim to keep inflation equal to its long-run or target rate. Since monetary policy operates by affecting aggregate demand, the way to move the price level back down towards the target path would be to raise interest rates in order to reduce aggregate demand. Since no such reduction would be necessary under an IT regime, the conditional variability of output would be lower.

4. The existence of imperfectly indexed long-term nominal contracts has implications for the effects of price-level shocks on the distribution of wealth under PLT and IT. This is an active area of research. See, for example, Doepke and Schneider (2006).
5. Some recent work analyzes the welfare benefits from reduced uncertainty surrounding the real value of the payoffs of nominal contracts. These studies take the existence of long-term nominal contracts as given. See, for example, Doepke and Schneider (2006).

In summary, the traditional view sees PLT as a trade-off between the longer-run benefits of increased price-level predictability and the short-run costs of increased variability of both prices and output. The contribution of the more recent literature on PLT has been to show that, under certain circumstances, PLT can actually lead to an improved trade-off between inflation and output variability.[6] Much of the focus of recent papers has been to investigate just how wide the range of these circumstances is.

3 A Free Lunch from Price-Level Targeting

Formal models in the early 1990s largely confirmed the traditional view concerning increased short-run variability of prices and output under PLT. Examples include Lebow, Roberts, and Stockton (1992), and Haldane and Salmon (1995). These models are based on adaptive expectations concerning future inflation, so that a change in monetary policy regime does not affect the way inflation expectations are formed.[7]

Svensson's (1999) seminal paper was the first to construct a model in which an improved short-run trade-off between output and inflation variability is possible. Before discussing the reasons underlying Svensson's result, it is important to note that in his paper, as in many of the papers discussed in this survey, society's preferences can be expressed in terms of a quadratic loss function that depends on variations in inflation and in the output gap:

$$\mathcal{L} = E_t \sum_{i=0}^{\infty} \left\{ \gamma {x_{t+i}}^2 + {\pi_{t+i}}^2 \right\}, \tag{1}$$

where x_t is the output gap at time t, which measures the proportional difference between output and the level of output that would prevail under complete price flexibility; π_t is the deviation at time t of the inflation rate from its long-run value; E_t is the expectations operator conditional on information available at time t; and $\gamma > 0$ is a positive parameter that measures the relative importance of deviations in inflation compared to deviations in output. Many papers simply posit this loss function. It can also be derived as an approximation of a social welfare function that depends on the expected utility of a

6. Most of this recent literature takes a longer-term perspective and examines the trade-off between the **unconditional** variances of inflation and output, which do not depend on the current state of the economy.
7. As discussed below, more recent models have reintroduced backward-looking expectations in order to help explain the observed degree of inflation persistence. It should not be surprising that, as the importance of backward-looking expectations increases in these models, the advantages of PLT tend to decrease.

representative household; see Woodford (2003) for details.[8]

Svensson assumes an aggregate supply curve of the following form:

$$x_t = \delta x_{t-1} + \alpha \left(\pi_t - E_{t-1}\pi_t\right) + \varepsilon_t,$$

where ε_t is a random supply shock and $\alpha > 0$. This equation can be rewritten as follows:

$$\pi_t = E_{t-1}\pi_t + \psi \left(x_t - \delta x_{t-1}\right) + \mu_t, \tag{2}$$

where $\psi \equiv 1/\alpha$ and $\mu_t \equiv -\varepsilon_t/\alpha$. Rewritten this way, the equation has the interpretation of a New Classical Phillips curve (McCallum 1994).

The central bank minimizes the loss function (1) subject to (2). By assumption, it cannot precommit to its future policies; for this reason, it reoptimizes in every period and its problem is a static one.[9] The central bank can observe the current value of the supply shock μ_t and can choose the inflation rate exactly.[10] The optimality conditions for this problem lead to a policy rule in which inflation depends on the current value of the output gap.

Given this solution, it is possible to solve for the unconditional variances of the inflation rate and the output gap. Both of these solutions will depend on γ, the relative weight attached to the output gap in (1). As the value of γ decreases, the central bank (and society) attaches less importance to fluctuations in output: it can be shown that the variance of output increases and the variance of inflation decreases, leading to a negative trade-off between the two variances that depends on γ.

It is also possible to solve for the central bank's optimal policy subject to a loss function that depends on the output gap and deviations of the **price level** from a target path. Such

8. For purposes of exposition, the natural level of output is assumed to be equal to its socially efficient level. Svensson (1999) assumes that the natural level of output is inefficiently low, so that the central bank is tempted to generate unexpectedly high inflation in order to boost output: in equilibrium, individuals rationally anticipate this temptation, and output is no greater, on average, than its natural level, but there is a positive bias to the inflation rate. Dittmar, Gavin, and Kydland (1999) show that this assumption is not required for the free lunch result.
9. If it can precommit to its future policies, it can, in general, attain a higher expected level of economic welfare. The distinction between optimal monetary policy under commitment and without it is crucial in this literature. Optimal policy under commitment is discussed in section 4.1.
10. In the real world, central banks affect inflation by affecting aggregate demand via their control over short-term nominal interest rates. In this simple model, the central bank observes all shocks prior to setting its interest rate, and the interest rate has an immediate effect on aggregate demand. An aggregate-demand equation can be added to the model, but it serves only to back out the interest rate that is required for the central bank to achieve its chosen inflation rate.

a loss function can be written as follows:

$$\mathcal{L}_p = E_t \sum_{i=0}^{\infty} \left\{ \gamma_p {x_{t+i}}^2 + (p_{t+i} - p^*_{t+i})^2 \right\}, \tag{3}$$

where we have added a subscript to the relative weight attached to deviations of the output gap in order to emphasize that the loss function is not the same as (1); p_{t+i} is the targeted price level; and p^*_{t+i} is the target path for the price level, which may be either constant or growing at a constant rate. Note that, in general, it is not possible to derive this loss function as an approximation of the true social welfare function. In contrast, (1) can be so derived: inflation has a direct impact on economic welfare because it influences the dispersion of prices across different firms and decreases the efficiency of production.[11] Howitt (2001) calls this instructing the central bank to act like a "Zen archer" by aiming at a target that is not society's true target.[12]

In this case, the central bank's optimality conditions give the price level as a function of the current value of the output gap, which means that inflation depends on the first difference of the output gap, rather than on the output gap itself. Again, it is possible to calculate the trade-off between the unconditional variance of inflation and the unconditional variance of the output gap as a function of γ_p. If the persistence of output as measured by the δ parameter in (2) is sufficiently high, the trade-off is unambiguously better[13] with a price-level target than with an inflation target.

One way of understanding this result is to note that, as the δ parameter increases, fluctuations in the output gap become more persistent. As δ approaches one, the output gap takes on the character of a random walk. With (1) as the central bank's objective function, inflation depends on the output gap, so that it, too, increasingly resembles a random walk as the persistence of output fluctuations increases. The optimal policy then entails persistent deviations of the inflation rate from the target rate, which increases the variability of inflation. With (3) as the objective function, the inflation rate remains stationary even when the output gap tends towards a random walk. With this objective function, the central bank worries about (and eliminates) the cumulative price-level errors

11. See Ambler (2007) for a detailed explanation.
12. Assigning an objective function other than the true social welfare function to the central bank has a long tradition in macroeconomics. One of the best known examples is Rogoff (1985), who constructs a model in which appointing a "conservative" central banker, who is more concerned than society as a whole with fighting inflation, could lead to an unambiguously better outcome, with lower inflation and the same average level of output.
13. That is to say, the variance of inflation is lower for a given value of the variance of the output gap. Equivalently, the variance of the output gap is lower for a given value of the variance of inflation.

that would arise when using (1) as the objective.

Another way of understanding Svensson's result is to consider that inflation expectations in his model are indirectly forward looking. In the presence of endogenous output-gap persistence, the central bank can affect the future trade-off between inflation and output variability by affecting the current output gap. As the output gap becomes more persistent, the central bank's ability to affect the future trade-off is enhanced. It can be shown that, if the output persistence is purely exogenous (arising from, for example, persistence in the error term μ_t), the central bank cannot affect the future trade-off between inflation and output variability, and there are no advantages to be had by assigning a Zen objective function to the central bank.

The importance of forward-looking expectations, either direct or indirect, was highlighted in a recent article by Cover and Pecorino (2005). They use the same basic model as Svensson (1999) and Dittmar, Gavin, and Kydland (1999), but change the assumption of the timing of the central bank's decisions. Cover and Pecorino suppose that the central bank must choose its optimal policy before knowing the current value of aggregate disturbances such as the μ_t shock in (2) above. In such a context, the aggregate-demand side of the economy plays an active role in the determination of macroeconomic equilibrium, rather than just recursively determining the nominal interest rate necessary to attain the central bank's chosen rate of inflation. In Cover and Pecorino's model, aggregate demand depends on the ex ante real interest rate, equal to the nominal interest rate minus expected inflation based on current information. Cover and Pecorino's main finding is that PLT is stabilizing (improves the trade-off between output and inflation variability) even with no endogenous output persistence. When there is a positive inflation shock under PLT, expected future inflation declines, which yields a higher real interest rate for any given level of the nominal interest rate. This reduces aggregate demand, which in turn reduces the equilibrium inflation rate in the current period.

The importance of forward-looking expectations is made even more clear when the advantages of PLT are considered in the context of New Keynesian models, whereby the New Classical Phillips curve (2) is replaced by a New Keynesian Phillips curve in which current inflation depends on expectations of future inflation based on current information. Results based on these models are discussed in detail in section 4, but mention can be made here of Dittmar and Gavin (2000), who use a modified version of (2), in which the only change is to replace the lagged expectation of current inflation with the current expectation of future inflation. They show that the trade-off between output and inflation variability

improves with an objective function that penalizes price-level deviations, irrespective of the degree of persistence of output fluctuations as measured by the δ parameter in (2).

In all these models, the optimal feedback rule for the central bank with the (3) objective function gives a relationship between the price level and the output gap, implying a relationship between inflation and the **change** in the output gap. The dependence of inflation on the lagged output gap introduces an element of **history dependence**. History dependence is one of the characteristics of optimal policy under commitment, as discussed in section 4.

4 The Robustness of the Free Lunch

4.1 The optimality of price-level stationarity under commitment in New Keynesian models

Most of the analyses of the relative benefits of PLT versus IT have used New Keynesian macroeconomic models, rather than models with the New Classical Phillips curve of Svensson's (1999) paper. These models have become workhorses for monetary policy analysis by both central banks and academic economists.[14]

The New Keynesian model is based on monopolistically competitive firms that produce goods that are imperfect substitutes. Firms set prices optimally, but they are unable by assumption to reoptimize their price in each period. When able to, they set a price that depends on their marginal cost of production and on their expectations for the overall price level over the period in which their price is expected to remain fixed. Under certain restrictions,[15] their price-setting decisions can be aggregated together to yield the basic New Keynesian Phillips curve, which can be written as follows:

$$\pi_t = \beta E_t \pi_{t+1} + \psi x_t + \mu_t, \tag{4}$$

where $0 < \beta < 1$ measures the subjective discount rate of the representative household, x_t measures the output gap (the proportional difference between output and its level under complete price flexibility), and $\psi > 0$ is a parameter that depends on underlying structural

14. The main model currently in use for internal forecasting purposes at the Bank of Canada, ToTEM, is an elaborate version of a New Keynesian model; see Murchison and Rennison (2006). For an exposition of the basic model and its application to monetary policy analysis, see Clarida, Galí, and Gertler (1999).
15. See Clarida, Galí, and Gertler (1999) or Ambler (2007) for details.

parameters.[16] The cost-push shock μ_t is generally appended to the equation in order to generate a meaningful trade-off between output and inflation.[17] It is possible to provide microfoundations for the cost-push shock by positing exogenous fluctuations in firms' demand elasticities and/or exogenous fluctuations in tax rates (Steinsson 2003).

Given the basic New Keynesian Phillips curve and a loss function of the form of (1), it is possible to solve for the central bank's optimal monetary policy problem, under the assumption that it can commit to its future policies. This assumption means that the central bank can use announcements of future policy to influence private agents' current expectations. Its ability to precommit to its future policies allows the central bank to attain a higher level of social welfare than otherwise. If the central bank can observe the current value of all aggregate disturbances when optimizing, it can directly choose the inflation rate to minimize the loss function (1), subject to (4). As in the model of Svensson (1999),[18] an aggregate-demand equation can be added to the model, but it serves only to calculate the short-term nominal interest rate necessary to meet the central bank's inflation target.

For a central bank that optimizes in period t, the bank's optimal rule for inflation has the following form:

$$\pi_t = -\lambda x_t,$$

$$\pi_{t+i} = -\lambda \left(x_{t+i} - x_{t+i-1} \right), \quad i > 0.$$

This solution has several interesting features. First, the central bank's choice of inflation at time t, when it optimizes, is different from its rule for choosing inflation in all subsequent periods. This difference is the source of the central bank's time-inconsistency problem. The central bank must be able to precommit credibly to its policy rule. If it was allowed to reoptimize at a later date, say $t + k$, it would choose $\pi_{t+k} = -\lambda x_{t+k}$, rather than $\pi_{t+k} = -\lambda(x_{t+k} - x_{t+k-1})$. If individuals expected the central bank to reoptimize, its announced policy would not be credible.

Second, for all periods after the initial period, the central bank's optimal policy is **history dependent**, since it depends on previous economic conditions as well as current economic

16. These include the degree of substitutability across the different types of goods produced by the monopolistically competitive firms, and the parameters of the firms' production functions.
17. Without the cost-push shock, the central bank can perfectly stabilize both inflation and the output gap. In New Keynesian models without cost-push shocks, complete stabilization of the price level is optimal. See King and Wolman (1999), Goodfriend and King (2001), and Goodfriend (2002) for a detailed explanation.
18. See footnote 10.

conditions. The history dependence of the optimal policy is a by-product of the central bank's ability to influence the expectations of the private sector. The private sector anticipates that future policy will be different because of changes in current conditions. In order for this to be credible, the central bank's current policies must depend on past conditions.

Third, the central bank's rule for inflation after period t is qualitatively of the same form as the optimal policy under discretion in the model of Svensson (1999). Inflation depends on the change in the output gap, rather than the level of the output gap. This feature of the central bank's optimal policy suggests that it may be possible for a central bank to achieve a more favourable trade-off by acting as a Zen archer when it cannot credibly commit to its future policies. This is Svensson's (1999) result, and, indeed, in some cases the same expected level of welfare can be attained under discretion as under commitment, as discussed in section 4.3.

Fourth, an important implication of this solution for optimal policy is that the price level is **stationary**. This result was first demonstrated by Woodford (1999) and by Clarida, Galí, and Gertler (1999). In response to a positive cost-push shock, inflation is initially positive, but less than the value of the cost-push shock itself as the central bank reduces aggregate demand in order to bring down inflation. Starting with the first period after the shock dissipates, inflation becomes negative, and the price level is gradually brought back to its initial pre-shock value.

It is easy to see why committing to reducing inflation in the future is beneficial. By committing to a reduction in future inflation, even after the shock has passed, current expectations of future inflation are reduced. According to the New Keynesian Phillips curve (4), current inflation depends directly on future expected inflation via its effects on the price-setting behaviour of firms that can change their prices in the current period. If these firms expect lower inflation in the future because of the central bank's credible commitment, they can set a lower current price and still be able to maintain the same relative price compared to competing firms in the future. For this reason, the central bank does not have to decrease aggregate demand as much in order to obtain a given reduction in current inflation. In other words, the trade-off between inflation and output in the current period improves, reducing the output loss associated with fighting inflation in the face of a positive cost-push shock. This, in turn, reduces inflation persistence, thereby reducing inflation variability.

While it is clear why committing to a reduction in future inflation favourably affects the output-inflation trade-off, it is not intuitively obvious why the optimal policy involves **completely** offsetting the initial increase in the price level. As shown in section 4.2, this result is not robust to the introduction of backward-looking elements in the New Keynesian Phillips curve.

4.2 When is price-level drift optimal?

One shortcoming of the standard New Keynesian Phillips curve is that it is unable to generate persistent inflation, as first pointed out by Fuhrer and Moore (1995). The typical response to this empirical shortcoming has been to add lagged inflation to the New Keynesian Phillips curve equation, which yields the so-called hybrid New Keynesian Phillips curve, which is of the form:

$$\pi_t = \chi E_t \pi_{t+1} + (1 - \chi)\, \pi_{t-1} + \psi x_t + \mu_t. \tag{5}$$

The usual justification for the presence of lagged inflation is that a fraction of firms are rule-of-thumb price-setters, setting their price based on past inflation, rather than on their rational expectation of future inflation (see, for example, Galí and Gertler 1999).

Steinsson (2003) generalizes the rule of thumb used by Galí and Gertler (1999). He supposes that rule-of-thumb price-setters set a price equal to the mean level of prices in the previous period adjusted for lagged inflation, and also adjusted to vary directly with the lagged output gap. He derives a modified Phillips curve that can be written as follows:

$$\pi_t = \chi_f \beta E_t \pi_{t+1} + \chi_b \pi_{t-1} + \psi_1 x_t + + \psi_2 x_{t-1} + \mu_t. \tag{6}$$

The relative weight on expected future inflation versus past inflation in this equation depends negatively on the fraction of rule-of-thumb price-setters in the economy.

Steinsson sets up and solves the central bank's optimal monetary policy problem under commitment. He also derives the central bank's loss function as a quadratic approximation of a representative household's utility function. Because of the presence of rule-of-thumb price-setters, the loss function is more complicated than (1), and depends on the change in the inflation rate, the lagged value of the output gap, and an interaction term between the lagged output gap and the change in inflation, in addition to current inflation and the

current value of the output gap.[19]

Steinsson shows that, as one would expect, with no rule-of-thumb price-setters in the economy, it is optimal to perfectly offset cost-push shocks, so that under the optimal monetary policy the price level is stationary. More importantly, he shows that, as the fraction of rule-of-thumb price-setters increases, the optimal degree of price-level offset decreases.

Why is it not optimal to eliminate price-level drift when expectations are not forward looking? An increase in the price level in New Keynesian models arises because those firms that are able to modify their output price choose to increase it. This creates a distortion in relative prices that reduces the efficiency of production.[20] If the central bank tries to bring the price level back to its initial level or path, firms whose relative prices are out of equilibrium may not be able to change their prices, and firms whose prices are on the equilibrium path may be pushed out of equilibrium. Minford (2004) puts it this way:

> The best thing to do strictly depends on the chances of being allowed to change your price. If it is low (the usual assumption), then it is best to keep the new price level as there is a low chance of those who already changed their price being allowed to change it back. If it is high (over 50%), then reversal could be worthwhile as there is a good chance that those who already changed could change back. The break-even chance is 50%; below this it is optimal to keep the new price level.

This merely exacerbates relative price distortions. To the extent that expectations are backward looking, the benefits in the short run from an improved trade-off between output and inflation are smaller, and it becomes optimal not to completely offset the initial shock to the price level, since fewer additional distortions are created.

It would be tempting to draw a general conclusion from Steinsson's (2003) paper and other papers with rule-of-thumb price-setters that, to the extent that price expectations are predetermined, price-level drift becomes optimal and the advantages of PLT diminish. However, such generalizations turn out to depend on the exact details of firms' price-setting

19. Woodford (2003) was the first to show how the introduction of rule-of-thumb price-setters introduces lagged terms (in his case, lagged inflation) into the quadratic approximation of the representative household's welfare function.
20. See Ambler (2007) for a detailed explanation.

behaviour. Ball, Mankiw, and Reis (2005) set up a model in which all firms face costs to update the information that they use when setting prices. They suppose that all firms can change their prices in each period, but only a fraction of firms receives information concerning the complete state of the economy. They show that the model generates inflation that is as persistent as the inflation generated by New Keynesian models with rule-of-thumb price-setters, and that optimal monetary policy (under commitment) implies a stationary price level, as in New Keynesian models with completely forward-looking price-setters. The benefits of PLT in Ball, Mankiw, and Reis's model derive from smaller prediction errors for firms setting a path for their prices.

The crucial difference between Ball, Mankiw, and Reis's model and most of the New Keynesian models with rule-of-thumb price-setters is that, given their limited information sets, firms' expectations are rational, taking into account both the structure of the economy and a knowledge of how monetary policy is determined. This provides a channel through which a credible commitment by the central bank to its future policy can modify the current trade-off between inflation and output variability.

4.3 Price-level targeting as a commitment device

Section 4.1 discussed the result that, in the absence of rule-of-thumb price-setters, price-level stationarity is optimal when the central bank can commit to its future policies. This result, along with results obtained using a New Classical Phillips curve by Svensson (1999) and others, suggests that assigning a loss function defined in terms of price-level deviations rather than inflation may allow central banks to move closer to the commitment solution even when they cannot precommit.

Vestin (2006) demonstrates an even stronger result. He uses a New Keynesian model with forward-looking price-setters and with a central bank that optimizes under discretion. He shows that, with no persistence in the cost-push shock, by assigning a loss function to the central bank that depends on price-level deviations, rather than inflation, and by choosing an appropriate weight on deviations in the output gap, the same level of social welfare can be achieved as with the optimal monetary policy under commitment using a quadratic approximation of the true social welfare function.

This is a remarkable result. It is well known that the level of social welfare that can be attained under commitment is necessarily at least as high as under discretion. Only in very special models and under special circumstances is this inequality not strict. The standard

New Keynesian model with forward-looking price-setters is one such case, but the result depends on assigning an objective function to the central bank that is different from the true social welfare function.[21]

When the cost-push shock in Vestin's model is persistent, it is no longer possible to replicate the commitment solution with discretionary monetary policy and a price-level target. However, assigning a price-level target to the central bank can still lead to an improvement in social welfare compared to the optimal discretionary monetary policy with an inflation target.

The intuition for Vestin's result is straightforward. Assigning the central bank an objective function that depends on price-level deviations, rather than inflation, has the effect of conditioning the expectations of agents in the private sector. A positive inflation shock due to a cost-push shock reduces expectations of future inflation. This has the same effect as if the central bank acted optimally and could commit to its future policy. Giving this objective function to the central bank is a substitute for commitment.

4.4 Average inflation targeting

Section 4.2 showed that the introduction of backward-looking rule-of-thumb price-setters implies that optimal monetary policy under commitment involves a certain amount of price-level drift in response to cost-push shocks. The amount of drift that is optimal increases as the fraction of rule-of-thumb price-setters increases.

A straightforward way to vary the amount of price-level drift under discretionary monetary policy is by targeting a moving average of current and past inflation rates, rather than the current inflation rate. By increasing the size of the window used to calculate the moving average, the amount of price-level drift in the long run in response to an unanticipated change in the price level is reduced. As the size of the window tends towards infinity, price-level drift is eliminated completely and the price level becomes stationary.[22]

21. This type of result has been criticized as being schizophrenic: even if the central bank is unable to precommit to an announced path for monetary policy, it must be able to precommit to its Zen target.

22. A potential side benefit of targeting a moving average of inflation is that it could make the task of communicating with the public simpler. Under PLT, in response to a positive inflation surprise, it would be necessary to revise downward the target inflation rate in order to get the price level to return to its growth path. With average inflation targeting, while it is true that the one-period inflation rate would have to be below the targeted average inflation rate if the average was above the target, as long as the central bank communicates in terms of the average inflation rate, rather than the period-by-period inflation rate, this should pose no special communication challenges. Issues related to the central bank's communication of its policy by the central bank are discussed in more detail in section 5.

Nessén and Vestin (2005) show that, under discretion, targeting average inflation can, under some circumstances, yield a superior outcome to both IT and PLT. Pure PLT dominates in a completely forward-looking model: this is not surprising, since Vestin (2006) shows that PLT can reproduce the optimum under commitment. As noted earlier, the optimal amount of price-level drift depends directly on the fraction of price-setters who use rule-of-thumb behaviour. Targeting average inflation allows the central bank to achieve this automatically: decreasing the size of the window used for calculating average inflation[23] increases the amount of price-level drift in the long run. As long as the fraction of rule-of-thumb price-setters is not too large, by choosing the optimal window size the central bank can do better than with pure IT or pure PLT. In some cases, the performance of average inflation targeting is very close to the optimal monetary policy under commitment. However, if the fraction of rule-of-thumb price-setters becomes too large, IT is better for economic welfare than targeting average inflation.

Nessén and Vestin also show that, when price-setting is dominated by rule-of-thumb, backward-looking firms, minimization of the true social welfare function under discretion dominates both PLT and average inflation targeting. This result is compatible with the intuition developed in section 4.2. When price-setting is dominated by rule-of-thumb price-setters, offsetting unexpected changes in the price level due to cost-push shocks merely creates additional distortions in relative prices, and yields no improvement in the short-run trade-off between output and inflation.

Nessén and Vestin's results on average inflation targeting are closely related to papers on so-called hybrid targeting (Batini and Yates 2003; Cecchetti and Kim 2005). In those papers, the central bank's loss function is made to depend on a weighted average of price-level deviations and inflation deviations. A positive weight on price-level deviations means no price-level drift in the very long run, but varying the relative weights on price-level deviations and inflation deviations changes the speed at which the price level is brought back to its target path. The behaviour of inflation and prices in the short and medium runs can be made to be very similar to their behaviour under average inflation targeting. The relative weights that yield the highest welfare depend in a complicated way on the parameters of the model. For some parameter values, hybrid targeting can dominate both IT and PLT. As in the case of average inflation targeting, this tends to occur in cases where price-setting is dominated by neither forward-looking nor rule-of-thumb price-setters.

23. The window size refers to the number of terms used to calculate the average.

5 Other Issues Related to Price-Level Targeting

5.1 Price-level targeting and the zero bound

The research program announced by the Bank of Canada in November 2006 proposes to look at both a lower inflation target and the potential advantages of PLT. The two sets of questions are closely related. A commonly stated objection to a lower inflation target is that it raises the possibility that nominal short-term interest rates will hit the so-called zero bound: the central bank cannot lower its target rate below zero given the availability of an alternative asset – namely, money balances – that always pays a zero nominal rate of interest. In response to large negative inflation shocks that call for expansionary monetary policy, the zero lower bound may become a binding constraint on monetary policy.

Some researchers have suggested that, for a given target inflation rate, adopting a PLT regime with price-level path that gives the same rate of inflation in the long run can help to avoid hitting the zero lower bound. The argument for why this would be the case is straightforward. A negative inflation shock under PLT is, if the regime is credible, expected to be followed by inflation that is higher than average in order to bring the price level back to its predetermined path. The channel through which monetary policy has real effects operates through the ex ante real interest rate. With expected inflation increasing in response to a negative inflation shock, the bank's target rate has to be reduced by less to achieve the same change in the real interest rate, compared to a situation in which inflation expectations remain approximately constant. For this reason, monetary policy has more leverage at or near the zero bound under PLT than under IT.

While the logic of this argument is very simple, a rigorous analysis of the effects of PLT on the zero-lower-bound problem are complicated by the inherent non-linearity of the effects of the lower bound. It acts as a constraint that binds only occasionally. As such, special mathematical techniques are required to simulate its economic impact in the context of dynamic stochastic general-equilibrium models. The most commonly used technique to numerically solve New Keynesian models involves linearizing the equilibrium conditions of the model in the neighbourhood of its long-run equilibrium. By construction, this technique is incapable of capturing the impact of the lower bound.

Two strategies are available. The first strategy, adopted in a recent paper by Eggertsson and Woodford (2003), is to set up a model that is simple enough to solve explicitly for the exact dynamic solution. Eggertsson and Woodford find that a simple PLT rule ameliorates

the zero-bound problem and approximates the true optimal monetary policy much more closely than a simple IT rule. The other option is to use a model that can only be solved numerically, and to use the appropriate numerical techniques to account for the effects of the zero bound. Wolman (2005) solves a dynamic general-equilibrium model using projection methods.[24] He also finds that simple rules that impose the stationarity of prices can help alleviate the lower-bound problem.

5.2 The effects of the targeting regime on contracting

Most of the literature comparing PLT and IT takes as given the type and degree of nominal rigidity across the two types of monetary policy regimes. It is important to note that the details of how prices are set in New Keynesian models is imposed by assumption. Any comparison between the two types of regime that holds the type of nominal rigidity constant is potentially vulnerable to the Lucas critique. Barnett and Engineer (2001, 132) note that:

> . . . 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.

This point is developed in a series of papers by Minford with various co-authors (Minford, Nowell, and Webb 2003; Minford and Peel 2003; Minford 2004). They build models with households that cannot insure against fluctuations in their real wage, and that have a strong interest in smoothing those fluctuations. The equilibrium degree of indexation of nominal wages to the price level is also endogenous, and can depend on the monetary policy regime that is in place. Minford and his various co-authors find that the optimal degree of wage indexation is lower under a PLT regime, and that this can lead to substantial welfare benefits. The superiority of PLT results from reducing fluctuations in the real wage in response to monetary shocks.

Amano, Ambler, and Ireland (2007) develop a model with nominal-wage rigidities and an endogenous degree of indexation to unexpected changes in the price level. They show, as in

24. See Aruoba, Fernández-Villaverde, and Rubio-Ramírez (2004) for a comparative survey of non-linear techniques for solving dynamic stochastic general-equilibrium models.

Minford's work with his co-authors, that the optimal degree of wage indexation is lower under a PLT regime. Improved welfare under PLT in their model comes from a different mechanism: it helps the economy respond better to real shocks, moving the labour market closer to Walrasian equilibrium.

Accounting for the effect of the monetary regime on contracting is difficult. The form of nominal rigidities that is built into New Keynesian models is taken as exogenous precisely because it is difficult to provide convincing and tractable foundations for these frictions. However, comparing social welfare across monetary policy regimes that are vulnerable to the Lucas critique can potentially give seriously misleading results. Endogenizing the degree of indexation and other features of price and wage setting across monetary policy regimes is an important and promising avenue for future research.

5.3 Prolonged movements in relative prices: which price level?

Most of the models that have been used to study the costs and benefits of PLT have contained either one or a small number of goods sectors. The models feature relative price changes across differentiated goods within a particular sector, which are always inefficient. The kinds of prolonged relative price swings across different broad classes of goods, such as commodities and manufactured goods, are absent from these models. Swings in volatile components of the CPI have led central banks such as the Bank of Canada to construct measures of "core" inflation that leave out those components. While the official target of the Bank of Canada remains the CPI, core inflation is tracked closely and used as one of many measures of the pressures on inflation over the short to medium term.

Ortega and Rebei (2006) address this issue in a multi-sector framework. They also analyze the relative advantages of PLT and IT, and of a weighted average of the two. They construct a small open-economy model of the Canadian economy with traded and non-traded sectors, and with nominal-price rigidities in both sectors (and differential pricing of traded goods between domestic and export markets), as well as nominal-wage rigidities. No clear advantages of PLT over IT emerge, and it is difficult to discern the key assumptions in their model that are responsible for their results. Aoki (2001) builds a somewhat simpler two-sector model. One of the sectors is a competitive, flexible-price sector, and one is a sticky-price sector with monopolistically competitive firms. Aoki finds that the optimal monetary policy in this framework entails the complete stabilization of inflation in the sticky-price sector alone. Insofar as relative prices must fluctuate in order to reduce fluctuations in the output gap, this allows prices in the flexible-price sector to do all

of the adjusting.

While Erceg, Henderson, and Levin (2000) do not focus on the choice of price index, their results are compelling. They build a forward-looking model with both nominal-wage and nominal-price rigidities. They show that it is optimal to target a weighted average of wage inflation and price inflation. The relative weight on wage inflation versus price inflation is related directly to the average length of nominal-wage rigidity compared to nominal-price rigidity. Their results are compatible with those of Aoki, and can be interpreted as a generalization of his results, since the relative degree of the rigidity of prices and wages is variable in their model.

This suggests that monetary policy should focus primarily (but not exclusively) on reducing fluctuations in prices that are relatively more rigid, while allowing more flexible prices to adjust relative to these rigid prices. This solution represents a compromise. It facilitates relative price adjustment across different broad categories of goods in the face of real shocks, while at the same time dampening inefficient relative price fluctuations across different monopolistic producers of the same category of good. Even though the Bank of Canada does not directly target core inflation, looking closely at a less volatile component of the overall price index is in keeping with the spirit of this result.

The result that past inflation surprises should not be offset is related to the discussion in this section. Even though most New Keynesian models have one homogeneous final good, price-setting is introduced via differentiated intermediate goods produced by monopolistically competitive firms. These firms are identical ex ante: they have identical production functions, are all affected in the same way by aggregate technology shocks, and their goods enter the aggregate production function for final output symmetrically. However, since price-setting is staggered (different firms set their prices at different times), these firms are not the same ex post. There are relative price differences across firms.[25] It is not generally optimal to induce firms that are currently setting their prices to lower them in order to compensate for unexpectedly high prices set by firms in previous periods.

5.4 The transition from inflation targeting to price-level targeting

Most formal comparisons of the welfare properties of the IT and PLT regimes are built on the premise that individuals understand perfectly the workings of both regimes so that

25. Price dispersion across firms is one of the main costs of inflation in New Keynesian models. See Ambler (2007) for a discussion.

their expectations are completely rational. These comparisons ignore the costs associated with a transition from an IT regime to a PLT regime, which would involve the private sector learning about the workings of the regime. The learning process itself could mean expectations that are more dispersed across individuals in the short run. The adjustment in expectations would present communication challenges to the central bank that is effecting the regime change.

If monetary policy announcements continued to be explained in terms of inflation, this would entail frequent revisions of targeted inflation rates in response to shocks that provoked deviations of the price level from its targeted path. This could lead to inflation expectations being less firmly anchored than under current IT regimes, even in the long run. For example, consensus forecasts of inflation over short and medium horizons in Canada seem to have coalesced around the Bank of Canada's targeted rate of inflation, which has been constant for over ten years. It would be possible to base communication concerning monetary policy on the price level itself, rather than the inflation rate. It is possible to imagine a situation in which an interest rate increase could be justified on the basis of the percentage deviation of the targeted price index from its targeted price path. It is not known what the effects of this would be on the expectations of a public that has been conditioned for a long time to think in terms of inflation rather than the level of prices.

Insofar as some degree of price-level drift in response to shocks is judged to be optimal (for example, because of the presence of rule-of-thumb price-setters), the central bank's problem of how to communicate its policy becomes potentially even more complicated. One possible way to simplify communication and to ease the costs of transition would be to target average inflation. Choosing the appropriate size of moving average could necessitate very little revision in the way the central bank communicates its policy decisions. All that would be required would be to redefine the targeted rate of inflation. Central banks that currently have explicit inflation targets are already implicitly using average inflation targeting. For example, the Bank of Canada tracks monthly data on the year-over-year rate of inflation, which is just the average of the twelve monthly inflation rates over the preceding year. Moving from a twelve-month average to an average defined using a different window size would likely entail minimal adjustment and learning by the public.

There are no formal models of the learning process during the transition from an IT regime to a PLT regime. There are, however, models of learning during the shift to a new monetary policy regime.[26] These could fruitfully be applied to the specific question of a

26. See Bullard (2006) for a survey.

transition from an IT regime to a PLT regime. The communication issue is less easily subject to formal modelling, but will have to be addressed by any central bank contemplating a move to PLT.

6 Conclusions

Table 1 summarizes the main results in the recent literature on PLT. The principal benefit from PLT results from the improved trade-off between output and inflation when expectations are forward looking, making it less costly for the central bank to reduce current inflation. The results of Ball, Mankiw, and Reis (2005) suggest that what is important is not the information set on which expectations are conditioned (i.e., whether individuals use all current information or only past information when forming expectations), but rather whether expectations pertain to current or future inflation, and whether expectations are formed using a knowledge of the model's structure. Even when price-setting is based on expectations of current inflation, as in the model of Svensson (1999), endogenous output persistence introduces an indirect channel through which the central bank can affect the trade-off between inflation and output. It is not necessary for price-setting itself to be based on expectations of future inflation. It suffices, as in the model of Cover and Pecorino (2005), for forward-looking expectations to affect the macroeconomic equilibrium.

Table 1
Summary of Main Results

In favour of PLT	
Forward-looking price-setters	Stationary prices under commitment
Commitment not possible	PLT a substitute for commitment
Information updating costly	Reduced forecast errors under PLT
Low trend inflation	Zero-bound problem less severe under PLT
Endogenous indexation	Improved response to real shocks under PLT
Against PLT	
Rule-of-thumb behaviour	Some price-level drift optimal

The benefits of PLT are not limited to this channel. Assigning a price-level target to a central bank that cannot commit to its optimal monetary policy can help it achieve superior outcomes. When information is costly, as in the model of Ball, Mankiw, and Reis (2005), PLT can be beneficial by reducing the average size of forecast errors. When trend inflation is low, PLT can help to alleviate zero-bound problems. Finally, when price and wage setting depend on the monetary policy regime, PLT can reduce the incentive for

contingent wage indexation and improve economic performance in the face of real shocks.

Only to the extent that expectations are not forward looking and do not take into account the model's structure (as is the case with rule-of-thumb price-setters in New Keynesian models) does PLT not help ameliorate the current trade-off between output and inflation. Undoing past inflation surprises creates additional distortions and is more costly in terms of output fluctuations. In general, it is preferable to let bygones be bygones. It is important to note that, in this context, backward-looking expectations mean expectations that are formed solely on the basis of the past realizations of variables, without taking into account the structure of the economy or a knowledge of the rules governing the central bank's conduct of monetary policy.

It will be necessary to explore the importance of backward-looking expectations in price-setting behaviour. The rule-of-thumb price-setting rules in current models are a convenient shortcut that helps generate the degree of inflation persistence that is observed in the data.[27] However, they are the least theoretically satisfactory and most ad hoc elements in modern New Keynesian models. It is unclear whether policy recommendations should be based on ad hoc modelling assumptions that are as vulnerable to the Lucas critique as are previous generations of macroeconomic models. One characteristic of the rule-of-thumb price-setting used in New Keynesian macroeconomic models is that it gives no weight whatsoever to monetary policy announcements. It should be possible to come up with price-setting rules that, while not fully compatible with rational expectations, take into account credible announcements of future monetary policy.

Insofar as backward-looking expectations remain an integral part of New Keynesian models (despite the lack of satisfactory microfoundations and despite their vulnerability to the Lucas critique), some amount of price-level drift in response to inflation surprises will be optimal. Section 4.4 showed that the amount of price-level drift in response to exogenous shocks can be varied by assigning to the central bank an objective function defined in terms of a moving average of past inflation rates as a target, rather than the current inflation rate. An important benefit of targeting average inflation is that it would considerably simplify a central bank's communication of its policy to the private sector. Couching its analysis in terms of average inflation rather than current inflation would minimize the changes in communication strategy from an IT regime to a PLT regime. As noted in

27. This may, in fact, be a false problem, to the extent that measured inflation persistence has been steadily decreasing in inflation-targeting countries. Perhaps inflation persistence in the past was related to the lack of credibility (and hence poorly conditioned inflation expectations) of previous monetary policy regimes.

section 5.4, central banks that track monthly data on year-over-year inflation rates are already implicitly targeting average inflation. It remains to be seen whether a moving-average inflation target would allow a central bank to attain under discretion the same level of social welfare as a central bank maximizing the true social welfare function under commitment, or at least to attain a level of social welfare that is very close to this optimum. In other words, it is an open question whether a result like Vestin's (2006) would apply in a model with backward-looking expectations. Even for a central bank able to precommit to its future policies, targeting a moving average of past inflation rates may be a straightforward and relatively transparent way of shifting at least partly towards a PLT regime from an IT regime. This should be one focus of future research.

The papers reviewed in this survey undertake normative analyses. They have to do with characterizing optimal monetary policy, and depend critically on whether the central bank is assumed to be able to commit to its future policies. This begs the question as to which assumption, discretion or commitment, is more appropriate as a positive description of central bank behaviour. This has been a controversial subject in the literature. Price levels in economies with IT regimes appear to have been non-stationary. This could be interpreted as evidence either of discretionary behaviour or of rule-of-thumb price-setting in the models used by the central banks to establish their policies.

References

Amano, R., S. Ambler, and P. Ireland. 2007. "Price-Level Targeting, Indexation, and Welfare." Bank of Canada. Forthcoming.

Ambler, S. 2007. "The Costs of Inflation in New Keynesian Models." Bank of Canada. Draft.

Aoki, K. 2001. "Optimal Monetary Policy Responses to Relative-Price Changes." *Journal of Monetary Economics* 48(1): 55–80.

Aruoba, S. B., J. Fernández-Villaverde, and R. F. Rubio-Ramírez. 2004. "Comparing Solution Methods for Dynamic Equilibrium Economies." Penn Institute for Economic Research Working Paper No. 04-003.

Ball, L., N. G. Mankiw, and R. Reis. 2005. "Monetary Policy for Inattentive Economies." *Journal of Monetary Economics* 52(4): 703–25.

Bank of Canada. 2006. *Renewal of the Inflation-Control Target: Background Information.* November.

Barnett, R. and M. Engineer. 2001. "When is Price-Level Targeting a Good Idea?" In *Price Stability and the Long-Run Target for Monetary Policy*, 101–43. Proceedings of a seminar held by the Bank of Canada, June 2000. Ottawa: Bank of Canada.

Batini, N. and T. Yates. 2003. "Hybrid Inflation and Price-Level Targeting." *Journal of Money, Credit and Banking* 35(3): 283–300.

Bullard, J. B. 2006. "The Learnability Criterion and Monetary Policy." *Federal Reserve Bank of St. Louis Review* 88(3): 203–17.

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

Cecchetti, S. G. and J. Kim. 2005. "Inflation Targeting, Price-Path Targeting, and Output Variability." In *The Inflation-Targeting Debate*, 173–95, edited by B. S. Bernanke and M. Woodford. Chicago: University of Chicago Press.

Clarida, R., J. Galí, and M. Gertler. 1999. "The Science of Monetary Policy: A New Keynesian Perspective." *Journal of Economic Literature* 37(4): 1661–1707.

Coulombe, S. 1998a. "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.

———. 1998b. "A Non-Paradoxical Interpretation of the Gibson Paradox." Bank of Canada Working Paper No. 98–22.

Cover, J. P. and P. Pecorino. 2005. "Price and Output Stability under Price-Level Targeting." *Southern Economic Journal* 72(1): 152–66.

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

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

Doepke, M. and M. Schneider. 2006. "Inflation and the Redistribution of Nominal Wealth." *Journal of Political Economy* 114(6): 1069–97.

Duguay, P. 1994. "Some Thoughts on Price Stability versus Zero Inflation." Bank of Canada. Draft.

Eggertsson, G. B. and M. Woodford. 2003. "The Zero Bound on Interest Rates and Optimal Monetary Policy." *Brookings Papers on Economic Activity* 1: 139–211.

Erceg, C. J., D. W. Henderson, and A. T. Levin. 2000. "Optimal Monetary Policy with Staggered Wages and Prices." *Journal of Monetary Economics* 46(2): 281–313.

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

Fuhrer, J. and G. Moore. 1995. "Inflation Persistence." *Quarterly Journal of Economics* 110(1): 200–223.

Galí, J. and M. Gertler. 1999. "Inflation Dynamics: A Structural Econometric Analysis." *Journal of Monetary Economics* 44(2): 195–222.

Goodfriend, M. 2002. "Monetary Policy in the New Neoclassical Synthesis: A Primer." *International Finance* 5(2): 165–91.

Goodfriend, M. and R. G. King. 2001. "The Case for Price Stability." In *Why Price Stability?* 53–94, edited by A. G. Herrero, V. Gaspar, L. Hoogduin, J. Morgan, and B. Winkler. Proceedings of the First ECB Central Banking Conference held in Frankfurt, Germany, November 2000. Frankfurt, European Central Bank.

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

Howitt, P. 2001. "Discussion of Michael Parkin's 'What Have We Learned About Price Stability?'" In *Price Stability and the Long-Run Target for Monetary Policy*, 260–65. Proceedings of a seminar held by the Bank of Canada, June 2000. Ottawa: Bank of Canada.

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

Lebow, D. E., J. M. Roberts, and D. J. Stockton. 1992. "Economic Performance under Price Stability." U.S. Board of Governors of the Federal Reserve Board, Economic Activity Section Working Paper No. 125.

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

Minford, P. 2004. "Monetary Policy — Should It Move onto a Price Level Target?" A. W. Phillips Memorial Lecture to the New Zealand Association of Economists, July.

Minford, P., E. Nowell, and B. Webb. 2003. "Nominal Contracting and Monetary Targets — Drifting into Indexation." *Economic Journal* 113(484): 65–100.

Minford, P. and D. Peel. 2003. "Optimal Monetary Policy: Is Price-Level Targeting the Next Step?" *Scottish Journal of Political Economy* 50(5): 650–67.

Murchison, S. and A. Rennison. 2006. *ToTEM: The Bank of Canada's New Quarterly Projection Model*. Technical Report No. 97. Ottawa: Bank of Canada.

Nessén, M. and D. Vestin. 2005. "Average Inflation Targeting." *Journal of Money, Credit and Banking* 37(5): 837–63.

Ortega, E. and N. Rebei. 2006. "The Welfare Implications of Inflation versus Price-Level Targeting in a Two-Sector, Small Open Economy." Bank of Canada Working Paper No. 2006-12.

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

Steinsson, J. 2003. "Optimal Monetary Policy in an Economy with Inflation Persistence." *Journal of Monetary Economics* 50(7): 1425–56.

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

Vestin, D. 2006. "Price-Level Targeting versus Inflation Targeting." *Journal of Monetary Economics* 53(7): 1361–76.

Wolman, A. L. 2005. "Real Implications of the Zero Bound on Nominal Interest Rates." *Journal of Money, Credit and Banking* 37(2): 273–96.

Woodford, M. 1999. "Optimal Monetary Policy Inertia." National Bureau of Economic Research Working Paper No. 7261.

———. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton: Princeton University Press.