doi: 10.1111/j.1467-6419.2009.00601.x

PRICE-LEVEL TARGETING AND STABILISATION POLICY: A SURVEY

Steve Ambler

CIRPÉE, Université du Québec à Montréal

Abstract. This paper surveys recent articles on the costs and benefits of price-level targeting, focusing its use as a tool for stabilisation policy. It discusses how price-level targeting can affect the short-run trade-off between output and inflation variability by influencing inflation expectations. It reviews how assigning an explicit price-level target to a central bank that is unable to commit to its future policies can improve economic performance. It surveys other potential benefits and costs. Among the costs, it underlines the importance of perfectly rational expectations for the optimality of price-level targeting, and an exacerbation of the time inconsistency problem.

Keywords. Inflation; Monetary policy; New Keynesian models; Price-level targeting

1. Introduction

There are currently 26 countries whose central banks formally use inflation targeting (IT) as their monetary framework (Lim, 2008). No country has abandoned IT except for Finland and Spain upon joining the euro zone. At least until the onset of the current worldwide recession, inflation and output fluctuations have been less volatile in IT countries. Whether or not IT itself was responsible for this 'great moderation',[1] it is clear that IT was adopted more by historical accident than because of a consensus that it was an optimal monetary policy framework. For example, in Canada's case, IT was adopted primarily to forestall the perceived threat of a wage–price spiral subsequent to the introduction of the federal goods and services tax (Laidler, 2007, p. 3). For this reason, it seems natural to ask whether there are other monetary policy frameworks that could be superior to IT.

One possible alternative to IT is price-level path targeting (PT).[2] Both IT and PT can allow for a positive inflation rate in the long run. The crucial distinction between the two regimes is how the central bank reacts to unexpected changes in inflation. Under IT, the central bank acts to bring inflation back to its target rate. It treats the effect of the inflation shock on the price level as a bygone. This means that a temporary shock to the inflation rate has a permanent effect on the price level. Under PT, the central bank acts to return the price level to its original targeted growth path.

Journal of Economic Surveys (2009) Vol. 23, No. 5, pp. 974–997

There is a substantial body of research that examines the costs and benefits of PT compared to IT.[3] This paper undertakes a survey of the modern literature on PT. Traditional analyses[4] focus on the long-run predictability of prices as the main benefit of PT, 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 PT can favourably affect the short-run trade-off between output and inflation variability by affecting expectations of future inflation. The seminal paper 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 loss function to the central bank with the price level as one of its arguments. Much of the recent literature on PT centres on analysing the robustness of Svensson's result in more complex economic environments than the one he studied. This paper aims to assess which conclusions are robust and which questions remain open.

The paper is structured as follows. Section 2 briefly reviews traditional arguments for and against PT. Section 3 takes up Svensson's free lunch argument. Section 4 looks at extensions to the environment of Svensson's paper. It discusses the result that in a standard New Keynesian model optimal monetary policy under commitment implies a stationary price level. It also reviews papers that show that PT is no longer optimal in New Keynesian models when some firms 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, it reviews papers that examine targeting a moving average of past inflation rates as a compromise between pure PT and pure IT. Section 5 examines other issues. It looks at the effects of a change in monetary regime on the incentive to index contracts, and how that affects the advantages of PT; the issue of how to deal with prolonged movements in relative prices; issues related to the transition between an IT regime and a PT regime; issues related to communication and transparency and the time inconsistency problem under PT. Section 6 attempts to draw some general conclusions and suggests possible avenues for future research.

2. Traditional Arguments For and Against Price-level Targeting

As noted in the introduction, under IT 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. Under PT, the conditional forecast-error variance of the price level remains finite at all time horizons.[5]

The long-run predictability of the price level under PT is the source of its intuitive appeal. It means that the real value of future payments specified contractually in nominal terms is more predictable than under IT. The reason for the prevalence of

long-term contracts that stipulate payments in nominal terms is not addressed here.[6] Under a PT regime, current prices convey intertemporal information, as argued by Coulombe (1998a, b).

If reduced price-level uncertainty is the main traditional argument in favour of PT, the traditional argument against PT is that it must raise the short-run variability of both inflation and output. The logic of this argument is straightforward. In response to a temporary, unexpected increase (decrease) in inflation in a PT 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. Formal models from the 1990s confirmed this traditional view: see Lebow *et al*. (1992) and Haldane and Salmon (1995).[7]

In summary, the traditional view sees PT 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 PT has been to show that, under certain conditions, PT can actually lead to an improved trade-off between inflation and output variability.[8] Much of the focus of recent papers has been to investigate just how wide the range of these conditions is.

3. A Free Lunch from Price-level Targeting

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. Svensson posits this loss function. In the New Keynesian models discussed in the next section, it can be derived as an

approximation to a social welfare function that depends on the expected utility of a representative household; see Woodford (2003) for details.[9]

Svensson assumes an aggregate supply curve of the following form:

$$x_t = \delta x_{t-1} + \alpha(\pi_t - E_{t-1}\pi_t) + \varepsilon_t \tag{2}$$

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

$$\pi_t = E_{t-1}\pi_t + \psi(x_t - \delta x_{t-1}) + \mu_t \tag{3}$$

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 minimises the loss function (1) subject to (3). By assumption, it cannot precommit to its future policies; for this reason, it reoptimises in every period and its problem is static. Optimal monetary policy in the absence of an ability to precommit to future policies is referred to in the literature as optimal discretionary monetary policy. The distinction commitment and discretion is crucial. Plosser (2007) gives good working definitions:[10]

> Commitment means delivering, in any particular situation, on past promises. In other words, the policymaker unequivocally will follow through on a promise made about future actions.
>
> Discretion, on the other hand, means that the policymaker is not bound by previous actions or plans and thus is free to make an independent decision every period.

If a central bank can commit it can, in general, attain a higher expected level of economic welfare. Optimal policy under commitment is discussed in Section 4.1. The central bank can observe the current value of the supply shock μ_t and can choose the inflation rate exactly.[11] 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 a loss function can be written as follows:

$$\mathcal{L}_p = E_t \sum_{i=0}^{\infty} \left\{ \gamma_p {x_{t+i}}^2 + (p_{t+i} - p^*_{t+i})^2 \right\} \tag{4}$$

where we have added a subscript to the relative weight attached to deviations of the output gap in order to emphasise 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.[12] Howitt (2001) notes that assigning a loss function of the form of equation (4) to the central bank is akin to getting it to act like a 'Zen archer' by aiming at a target that is not society's true target.[13]

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 (3) is sufficiently high, the trade-off is unambiguously better[14] 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 also 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 (4) 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 that would arise when using (1) as its objective.

Another interpretation of the result is that it alleviates a fundamental time inconsistency problem when output is persistent.[15] Consider the economy's response to a negative supply shock. Inflation increases and output falls in response to the shock. Because output fluctuations are persistent, the output gap is expected to remain negative for several periods. As long as the output gap remains negative, the central bank will be tempted to create inflation in order to narrow the output gap, to the extent that the marginal cost of additional inflation is less than the marginal benefit of reducing the output gap. Individuals will realise this, leading to higher inflation expectations. This in turn has a negative effect on output via equation (2). If the central bank maximises a loss function defined in terms of the price level, it will be less tempted to create inflation to reduce the output gap, the so-called 'stabilisation bias' will be reduced,[16] and output will be higher due to lower inflation expectations. The price-level target effectively substitutes for a commitment not to create inflation in the future.

Yet 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 paper by Cover and Pecorino (2005). They use the same basic model as Svensson (1999) and Dittmar *et al.* (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 (3) previously mentioned. 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 PT is stabilising (improves the trade-off between output and inflation variability) even with no endogenous output persistence. When there is a positive inflation shock under PT, 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 PT are considered in the context of New Keynesian models, in which the New Classical Phillips curve (3) 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 (3), 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 penalises price-level deviations, irrespective of the degree of persistence of output fluctuations as measured by the δ parameter in (3).

In all these models, the optimal feedback rule for the central bank with the (4) 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 studies of the relative benefits of PT 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. Clarida *et al.* (1999) give a good summary of the canonical model.

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 reoptimise their price in each period. The exogeneity of price stickiness in New Keynesian models is a potentially serious shortcoming. The canonical model also abstracts entirely from money. Woodford (2003, chapter 4) shows that the predictions of the canonical model are similar to those of a model in which money appears in individuals' utility functions, as long as utility functions are separable in real balances. 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,[17] their price-setting decisions can be aggregated together to yield the basic New Keynesian Phillips curve, which can be written as follows:

$$\pi_t = \beta E_t \pi_{t+1} + \psi x_t + \mu_t \tag{5}$$

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 parameters.[18] The cost-push shock μ_t is generally appended to the equation in order to generate a meaningful trade-off between output and inflation. Without the cost-push shock, the central bank can perfectly stabilise both inflation and the output gap. The price level should be kept perfectly constant. See King and Wolman (1999), Goodfriend and King (2001) and Goodfriend (2002) for a detailed explanation. 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 has the important consequence 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 optimising, it can directly choose the inflation rate to minimise the loss function (1), subject to (5). As in the model of Svensson (1999),[19] 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 optimises in period t, the bank's optimal rule for inflation has the following form:

$$\pi_t = -\lambda x_t \tag{6}$$

at time t, and

$$\pi_{t+i} = -\lambda(x_{t+i} - x_{t+i-1}) \quad i > 0 \tag{7}$$

where $\lambda > 0$ is a constant that depends on the structural parameters of the model. This solution has several interesting features. First, the central bank's choice of inflation at time t, when it optimises, 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 reoptimise 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 reoptimise, 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 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 *et al.* (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. The stationarity of the price level can be seen by noting that equation (7) is just the first difference of an equation relating the price level to the output gap. In the long run, the output gap is zero, and this determines the level of prices. Because the optimal policy implies price-level stationarity, the free lunch result extends to the canonical New Keynesian model. No output persistence in the form of a lagged output term in the Phillips curve is required for this result. In this sense, the free lunch result is

robust to replacing the New Classical Phillips curve by the New Keynesian Phillips curve.

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 (5), 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[20] 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

$$\pi_t = \chi E_t \pi_{t+1} + (1 - \chi)\pi_{t-1} + \psi x_t + \mu_t \tag{8}$$

where $0 \leq \chi \leq 1$. The usual justification for the presence of lagged inflation is that a fraction of firms are rule-of-thumb price-setters. They set their price based on past inflation, rather than on their rational expectation of future inflation.[21]

Steinsson (2003) generalises 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:

$$\pi_t = \chi_f \beta E_t \pi_{t+1} + \chi_b \pi_{t-1} + \psi_1 x_t + + \psi_2 x_{t-1} + \mu_t \tag{9}$$

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.[22]

Steinsson shows that (as one would expect) with no rule-of-thumb price-setters in the economy it is optimal to offset cost-push shocks completely, 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. 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.

To summarise, PT is not robust to rule-of-thumb price setting by firms. This is not surprising. It reconfirms the results of Lebow *et al.* (1992) and Haldane and Salmon (1995). They showed, in models with backward-looking expectations formation, that PT does not beat IT.

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 PT diminish. However, such generalisations turn out to depend on the exact details of firms' price-setting behaviour. Gaspar *et al.* (2007) introduce lagged inflation into the New Keynesian Phillips curve via partial indexation of prices to past inflation. The crucial distinction compared with rule-of-thumb behaviour is that firms that set prices do so in a forward-looking way, but firms can partially adjust their prices every period to past inflation, even in periods where they cannot reoptimise their price. In their model, it is optimal for the central bank to offset

shocks to the price level completely unless indexation is complete. Ball *et al.* (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 PT in Ball *et al.*'s model derive from smaller prediction errors for firms setting a path for their prices.

The crucial difference between both Ball *et al.* (2005) and Gaspar *et al.* (2007) on the one hand and New Keynesian models with rule-of-thumb price-setters on the other is that in the former firms' price-setting rules are forward-looking, and their expectations take 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. If the government assigns the target, this begs the question of the credibility of the government's commitment: this issue is beyond the scope of the current paper.

Vestin (2006) demonstrates an even stronger result. He uses a New Keynesian model with forward-looking price-setters and with a central bank that optimises 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.[23]

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.[24]

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 some price-level drift in response to cost-push shocks is optimal, even if commitment is possible. 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.[25]

Nessén and Vestin (2005) show that, under discretion, targeting average inflation can, under some circumstances, yield a superior outcome to both IT and PT. Pure PT dominates in a completely forward-looking model: this is not surprising, since Vestin (2006) shows that PT 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[26] 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 PT. 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, minimisation of the true social welfare function under discretion dominates both PT 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 PT. 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.

5. Other Issues Related to Price-level Targeting

5.1 *Price-level Targeting and the Zero Bound*

The research programme announced by the Bank of Canada (2006) included looking at both a lower inflation target and the potential advantages of PT. 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 more frequently. 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 PT 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 PT, if the regime is credible, is 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 PT than under IT. Because of this, it is less likely that the central bank's policy rate will be constrained by its effective lower bound. In this context, the historical experience of the gold standard is relevant. The gold standard can be interpreted as a form of PT in which the target was the price of one commodity rather than a broad price index. Short-term interest rates remained positive even through periods of overall price deflation during the gold standard era.

While the logic of this argument is simple, a rigorous analysis is 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 for solving New Keynesian models involves linearising the equilibrium conditions of the model in the neighbourhood of its steady state. By construction, this technique is incapable of capturing the impact of the lower bound.[27]

Two strategies are available. The first, followed by Eggertsson and Woodford (2003), is to set up a model that is simple enough to solve explicitly for the exact dynamic solution. They find that a simple PT rule ameliorates the zero-bound problem and approximates the true optimal monetary policy much more closely than a simple IT rule. The second is 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.[28] He also finds that simple rules that impose the stationarity of prices can help alleviate the lower-bound problem.

The possible advantages of PT close to the zero bound are of more than theoretical interest. Currently (June 2009), several major central banks have moved their policy rates close to zero and are actively seeking ways to make their monetary policies even more expansionary. PT has received some attention in this respect, for example from Mankiw (2008):

> A credible promise of subsequent price reversal after any deflation ensures that long-term expected inflation stays close to the inflation rate implied by the Fed's target price path. Monetary economists will recognise that this policy is price-level targeting rather than inflation targeting.

If inflation is expected to remain very low for some time, followed by a return to the targeted inflation rate (under IT), the average expected inflation rate over this period would be close to zero. Under a credible commitment to a price-level path, average expected inflation would be equal to the slope of the price-level path (the long-run inflation rate). For the same time path of short-term nominal interest rates, the long-term real interest rate would be lower by the difference in average expected inflation. A nominal interest rate stuck close to zero is therefore more expansionary under PT than under IT.

5.2 *The Effects of the Monetary Regime on Contracting*

Most of the literature comparing PT 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, p. 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 and Peel, 2003; Minford *et al.*, 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 PT regime, and that this can lead to substantial welfare benefits. The superiority of PT results from reducing fluctuations in the real wage in response to monetary shocks.

Amano *et al.* (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 Minford's work with his co-authors, that the optimal degree of wage indexation is lower under a PT regime. Improved welfare under PT 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. Endogenising 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*

Most of the models that have been used to study the costs and benefits of PT 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 consumer price index (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 analyse the relative advantages of PT 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 PT 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 the other is a sticky-price sector with monopolistically competitive firms. Aoki finds that the optimal monetary policy in this framework entails the complete stabilisation of inflation in the sticky-price sector alone. In so far 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 *et al.* (2000) do not focus on the choice of price index, their results are suggestive. 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 generalisation of his results, since the relative degree of the rigidity of prices and wages is variable in their model.

These results imply that PT using the full CPI will in general not be optimal. Monetary policy should focus primarily on reducing fluctuations in prices that are relatively more rigid, while allowing more flexible prices to adjust relative to these rigid prices.[29] 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.

5.4 *The Transition from Inflation Targeting to Price-level Targeting*

Most formal comparisons of the welfare properties of the IT and PT regimes are built on the premise that individuals understand perfectly the workings of both regimes so that their expectations are completely rational. These comparisons ignore the costs associated with a transition from an IT regime to a PT 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: this is discussed in the next section.

There has been some work on modelling learning during the shift to a new monetary policy regime.[30] Gaspar *et al.* (2007) built a model of adaptive learning applied to the transition to a PT regime. They conclude that learning reduces the gains to be had from switching to PT, but the net gains remain positive unless learning is implausibly slow.

5.5 *Communication*

Central banks under IT have communicated their forecasts in terms of inflation rates. Inflation expectations in countries with inflation-targeting central banks have coalesced around their targeted inflation rate. This would seem to suggest that IT has achieved a large measure of credibility in the eyes of the public where it is the official policy. 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. 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.

In so far 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. As shown in Section 4.4, the amount of price-level drift in response to exogenous shocks can be varied by assigning an average inflation objective to the central bank. 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 12 monthly inflation rates over the preceding year. Moving from a 12-month average to an average defined using a different window size would likely entail minimal adjustment and learning by the public.

5.6 *Time Inconsistency*

In order for PT to be successful, both rational expectations and a credible precommitment to its future policies on the part of the central bank are crucial. However, this commitment is bound up with a fundamental time inconsistency problem. In response to a positive inflation shock, a commitment to fighting future inflation lowers the cost of fighting current inflation because of the effects of the announcement of future policy on expectations. After the central bank reaps the benefits of this announcement, it is in its interests and the interests of society as a whole to renege on its announced policies. Bringing inflation below its long-run target rate in order to return the price level to its preannounced path is costly.

Table 1. Summary of Main Results.

Arguments in favour of PT	
Situation	**Advantages of PT**
Price-setters are forward-looking	PT is the optimal monetary policy under commitment
Commitment is not possible	PT can substitute for commitment
Updating information is costly	PT leads to reduced forecast errors
Trend inflation is low	The zero-bound problem is less severe under PT
Indexation is endogenous to the policy regime	PT improves the response of the economy to real shocks
Arguments against PT	
Situation	**Disadvantages of PT**
Expectations are not fully rational	Some price-level drift is optimal
Persistent relative price changes are required	Targeting the overall price level is not optimal
Learning about a new PT regime is required	PT can be costly if learning is slow
Communication strategies need to adapt to PT	Inflation expectations could become more volatile
Commitment necessary for PT to succeed	PT exacerbates the time inconsistency problem

In some situations, the time inconsistency problem can be particularly acute under PT. As argued in Section 5.1, if short-term nominal interest rates are stuck at or close to zero, a commitment to return prices to their target path will result in lower real interest rates, if credible. However, this involves a commitment down the road to engineer inflation that may be quite a bit higher than the long-run target rate for quite some time.

6. Conclusions

Table 1 summarises the main results from the recent literature on PT. Its principal benefit results from the improved trade-off between output and inflation when expectations are forward-looking. The results of Ball *et al.* (2005) suggest that what is important is not the information set used by individuals, but rather whether expectations pertain to current or future inflation, and whether they 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 suffices, as in Cover and Pecorino (2005), for forward-looking expectations to affect the macroeconomic equilibrium.

The benefits of PT 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 *et al.* (2005), PT can be beneficial by reducing the average size of forecast errors. When trend inflation is low, PT can help to alleviate zero-bound problems. Finally, when price and wage setting depend on the monetary policy regime, PT can reduce the incentive for contingent wage indexation and improve economic performance in the face of real shocks.

With expectations that do not take into account the model's structure, like the case of rule-of-thumb price-setters in New Keynesian models, PT does not improve 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, meaning that PT is not optimal.

It will be necessary to explore the importance of such expectations in price-setting behaviour. Rule-of-thumb price-setters are a convenient shortcut that helps generate inflation persistence, but they are also the least theoretically satisfactory and most *ad hoc* element 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. Rule-of-thumb price-setters give 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. The work by Ball *et al.* (2005) may point the way here: expectations in their model are forward-looking but (rationally) do not automatically incorporate the latest available information.

One context in which non-rational or adaptive expectations are likely to prevail is during a transition between monetary policy regimes. The results of Gaspar *et al.* (2007) are encouraging for PT, but there are many different ways to model learning and different possibilities should be explored before we judge that the advantages of PT are robust to learning.

The issues of the appropriate choice of a price index when changes in relative prices across sectors are necessary, of the appropriate communication strategy in a PT regime, and of how best to deal with time inconsistency will require further study before a switch to PT can be realistically contemplated. The issues are interrelated: stabilising a narrow price index could lead to volatile fluctuations in the overall CPI, making the central bank's problem of communicating its strategy more difficult and undermining the credibility of its commitment to announced future policies.

This survey has stressed the importance of commitment and how in some cases assigning an explicit price level target can substitute for commitment. It is not obvious to which institutional arrangements these abstract ideas correspond.[31] As noted in Section 5.6, time inconsistency may be particularly severe under PT.

Further study is needed concerning the institutional arrangements that would be needed to implement a PT regime and guarantee its credibility.

Finally, deciding on a switch to PT on the basis of the results from simulating New Keynesian models will always be a leap of faith. We will have to decide whether the intuitions that these models provide are robust enough to hold without the extreme abstractions of the New Keynesian approach: exogenous nominal rigidities and (at least in the basic models) no money or financial sector. Even abstracting from capital accumulation may lead us to put too much faith in the predictions of New Keynesian models concerning optimal monetary policy under discretion, if the problem of multiple equilibria is swept under the rug.

The papers reviewed in this survey are normative. They have to do with characterising 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.

Notes

1. Dotsey (2006) surveys the evidence.
2. The Bank of Canada (2006) announced a research programme to investigate the effects of lowering its inflation target below 2% and the benefits and costs of PT.
3. Because PT has never actually been implemented except for a short period in Sweden in the 1930s, this research relies on simulating calibrated or estimated model economies.
4. The older literature on PT extends back to Keynes, Fisher, Wicksell and others. See Duguay (1994) for a cogent survey.
5. In terms of time-series analysis, the price level is 'trend stationary' or I(0) under PT, and 'difference stationary' or I(1) under IT.
6. The existence of imperfectly indexed long-term nominal contracts has implications for the effects of price-level shocks on the distribution of wealth under PT and IT. This is an active area of research. See, for example, Doepke and Schneider (2006), Meh *et al.* (2008) and Meh and Terajima (2009).
7. These models were based on adaptive expectations concerning future inflation, so that a change in monetary policy regime did not affect the way inflation expectations are formed.
8. Most of this recent literature looks at the trade-off between the unconditional variances of inflation and output, which do not depend on the current state of the economy.
9. 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 *et al.* (1999) show that this assumption is not required for the free lunch result.

10. His speech also contains a clear discussion of the types of institutional arrangements that can lead to commitment in the real world.
11. 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.
12. See Ambler (2007) for a detailed discussion. Price dispersion is a feature of New Keynesian models, but not of the New Classical Phillips curve of the previous section.
13. 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.
14. That is, 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.
15. The author would like to thank an anonymous referee for pointing out this interpretation.
16. For a good intuitive explanation of stabilisation bias under the New Classical Phillips curve see Gärtner (2003).
17. See Clarida *et al.* (1999) or Ambler (2007) for details.
18. 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.
19. See footnote 11.
20. Benati (2008) presents evidence that this may be a red herring: after accounting for changes in monetary policy regimes, he finds little evidence of inflation persistence.
21. See for example Galí and Gertler (1999).
22. 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.
23. A potential problem with optimal policies under discretion, as pointed out by Blake and Kirsanova (2008), is that the equilibrium with optimal discretionary monetary policy may not be unique. Vestin's model abstracts from capital accumulation. Blake and Kirsanova's paper implies that this abstraction is crucial.
24. This type of result has been criticised 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.
25. 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 PT, 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 are discussed in more detail in Section 5.

26. The window size refers to the number of terms used to calculate the average.
27. All perturbation techniques that use smooth approximations around the steady state have the same shortcoming.
28. See Aruoba *et al.* (2004) for a comparative survey of non-linear techniques for solving dynamic stochastic general-equilibrium models.
29. Once again, the historical experience of the gold standard is relevant. The gold standard kept a flexible commodity price fixed while allowing all other prices to adjust to it, including prices subject to nominal rigidities. The studies reviewed in this section suggest that this would be highly suboptimal, explaining the difficult adjustment of Great Britain during the period between the two World Wars, when prices and nominal wages had to adjust downward after the price of gold was pegged at its level prior to the First World War.
30. See Bullard (2006) for a survey.
31. For surveys of different methods for achieving commitment and mechanisms that can substitute for commitment, see Gärtner (2003) and Plosser (2007).

References

Amano, R., Ambler, S. and Ireland, P. (2007) Price-level targeting, indexation, and welfare. Bank of Canada, Ottawa. Draft.

Ambler, S. (2007) The costs of inflation in New Keynesian models. *Bank of Canada Review* Winter: 5–14.

Aoki, K. (2001) Optimal monetary policy responses to relative-price changes. *Journal of Monetary Economics* 48: 55–80.

Aruoba, S.B., Fernández-Villaverde, J. and Rubio-Ramírez, R.F. (2004) Comparing solution methods for dynamic equilibrium economies. Penn Institute for Economic Research Working Paper No. 04-003.

Ball, L., Mankiw, N.G. and Reis, R. (2005) Monetary policy for inattentive economies. *Journal of Monetary Economics* 52: 703–725.

Bank of Canada (2006) *Renewal of the Inflation-control Target: Background Information*. Ottawa: Bank of Canada.

Barnett, R. and Engineer, M. (2001) When is price-level targeting a good idea? In *Price Stability and the Long-run Target for Monetary Policy* (pp. 101–143). Proceedings of a seminar held by the Bank of Canada, June 2000. Ottawa: Bank of Canada.

Batini, N. and Yates, T. (2003) Hybrid inflation and price-level targeting. *Journal of Money, Credit and Banking* 35: 283–300.

Benati, L. (2008) Investigating inflation persistence across monetary regimes. *Quarterly Journal of Economics* 123: 1005–1060.

Blake, A. and Kirsanova, T. (2008) Discretionary policy and multiple equilibria in LQ RE models. University of Exeter. Draft.

Bullard, J.B. (2006) The learnability criterion and monetary policy. *Federal Reserve Bank of St. Louis Review* 88: 203–217.

Calvo, G.A. (1983) Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics* 12: 383–398.

Cecchetti, S.G. and Kim, J. (2005) Inflation targeting, price-path targeting, and output variability. In B.S. Bernanke and M. Woodford (eds), *The Inflation-targeting Debate* (pp. 173–195). Chicago, IL: University of Chicago Press.

Coulombe, S. (1998a). The intertemporal nature of information conveyed by the price system. In *Price Stability, Inflation Targets, and Monetary Policy* (pp. 3–28). Proceedings of a conference held by the Bank of Canada, May 1997. Ottawa: Bank of Canada.

Coulombe, S. (1998b) A non-paradoxical interpretation of the Gibson paradox. Bank of Canada Working Paper No. 98-22.

Cover, J.P. and Pecorino, P. (2005) Price and output stability under price-level targeting. *Southern Economic Journal* 72: 152–166.

Dittmar, R. and Gavin, W. (2000) What do New-Keynesian Phillips curves imply for price-level targeting? *Federal Reserve Bank of St. Louis Review* 82: 21–30.

Dittmar, R., Gavin, W. and Kydland, F. (1999) The inflation–output variability tradeoff and price-level targets. *Federal Reserve Bank of St. Louis Review* 81: 23–31.

Doepke, M. and Schneider, M. (2006) Inflation and the redistribution of nominal wealth. *Journal of Political Economy* 114(6): 1069–1097.

Dotsey, M. (2006) A review of inflation targeting in developed countries. *Federal Reserve Bank of Philadelphia Business Review* Third Quarter: 10–20.

Duguay, P. (1994) Some thoughts on price stability versus zero inflation. Bank of Canada, Ottawa. Draft.

Eggertsson, G.B. and Woodford, M. (2003) The zero bound on interest rates and optimal monetary policy. *Brookings Papers on Economic Activity* 1: 139–211.

Erceg, C.J., Henderson, D.W. and Levin, A.T. (2000) Optimal monetary policy with staggered wages and prices. *Journal of Monetary Economics* 46: 281–313.

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

Galí, J. and Gertler, M. (1999) Inflation dynamics: a structural econometric analysis. *Journal of Monetary Economics* 44: 195–222.

Gärtner, M. (2003) Monetary policy and central bank behaviour. In C.K. Rowley and F. Schneider (eds), *The Encyclopedia of Public Choice*. Boston, MA: Kluwer Academic.

Gaspar, V., Smets, F. and Vestin, D. (2007) Is time ripe for price level path stability? European Central Bank Working Paper No. 818.

Goodfriend, M. (2002) Monetary policy in the new neoclassical synthesis: a primer. *International Finance* 5: 165–191.

Haldane, A.G. and Salmon, C.K. (1995) Three issues on inflation targets: some United Kingdom evidence. In A.G. Haldane (ed.), *Targeting Inflation* (pp. 170–201). 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* (pp. 260–265). Proceedings of a seminar held by the Bank of Canada, June 2000. Ottawa: Bank of Canada.

King, R.G. and Wolman, A.L. (1999) What should the monetary authority do when prices are sticky? In J.B. Taylor (ed.), *Monetary Policy Rules* (pp. 349–398). Chicago, IL: University of Chicago Press.

Laidler, D. (2007) Better late than never: towards a systematic review of Canada's monetary policy regime. C.D. Howe Institute Commentary No. 252, Toronto.

Lebow, D.E., Roberts, J.M. and Stockton, D.J. (1992) Economic performance under price stability. US Board of Governors of the Federal Reserve Board, Economic Activity Section Working Paper No. 125.

Lim, G. (2008) Inflation targeting. Available at http://insights.unimelb.edu.au/vol4/11_Lim.html last accessed 10 March (2009).

Mankiw, N.G. (2008) What is the fed to do? Available at http://gregmankiw.blogspot.com/2008/11/what-is-fed-to-do.html last accessed 10 March (2009).
McCallum, B.T. (1994) A semi-classical model of price-level adjustment. *Carnegie-Rochester Conference Series on Public Policy* 41: 251–284.
Meh, C. and Terajima, Y. (2009) Expected inflation and redistribution in Canada. *Bank of Canada Review* Spring: 43–50.
Meh, C., Ríos-Rull, J.-V. and Terajima, Y. (2008) Aggregate and welfare effects of redistribution of wealth under inflation and price-level targeting. Bank of Canada Working Paper No. 2008-31.
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. and Peel, D. (2003) Optimal monetary policy: is price-level targeting the next step? *Scottish Journal of Political Economy* 50: 650–667.
Minford, P., Nowell, E. and Webb, B. (2003) Nominal contracting and monetary targets – drifting into indexation. *Economic Journal* 113: 65–100.
Nessén, M. and Vestin, D. (2005) Average inflation targeting. *Journal of Money, Credit and Banking* 37: 837–863.
Ortega, E. and Rebei, N. (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.
Plosser, C. (2007) Credibility and commitment. Available at http://www.philadelphiafed.org/publications/speeches/plosser/2007/03-06-07_ny-assoc-bus-econ.cfm (last accessed 14 June 2009).
Rogoff, K. (1985) The optimal degree of commitment to an intermediate monetary target. *Quarterly Journal of Economics* 100: 1169–1189.
Steinsson, J. (2003) Optimal monetary policy in an economy with inflation persistence. *Journal of Monetary Economics* 50: 1425–1456.
Svensson, L.E.O. (1999) Price-level targeting versus inflation targeting: a free lunch? *Journal of Money, Credit, and Banking* 31: 277–295.
Vestin, D. (2006) Price-level targeting versus inflation targeting. *Journal of Monetary Economics* 53: 1361–1376.
Wolman, A.L. (2005) Real implications of the zero bound on nominal interest rates. *Journal of Money, Credit and Banking* 37: 273–296.
Woodford, M. (1999) Optimal monetary policy inertia. National Bureau of Economic Research Working Paper No. 7261.
Woodford, M. (2003) *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton, NJ: Princeton University Press.