Price-Level Targeting and Relative-Price Shocks

Stephen Murchison, Canadian Economic Analysis

- *Since 2006, the Bank of Canada has spearheaded a research program to examine the merits of switching to a regime that targets the price level rather than the rate of inflation.*
- *This article reviews model-based research focused on examining the relative merits of the two regimes in a small open economy, such as Canada's, that is susceptible to large and persistent shocks to its terms of trade. Research on the optimal price index under price-level targeting is also discussed.*
- *The balance of evidence suggests that the two regimes, implemented through simple policy rules, are quite similar in their ability to stabilize inflation, the output gap, and interest rates.*
- *Conditional on adopting price-level targeting, the overall CPI would represent close to an ideal index to target.*

In the autumn of 2006, researchers at the Bank of Canada embarked on an ambitious program to explore the potential welfare gains of switching from the Bank's current framework of targeting the *rate of change* in prices (i.e., inflation) to targeting the *price level*.[1] While research to date had suggested possible gains, several questions pertinent to Canada were identified as requiring further research. Among these was, *What are the relative merits of inflation targeting versus price-level targeting in an open economy susceptible to large and persistent terms-of-trade shocks?* (Bank of Canada 2006).

At issue is whether a central bank that targets an aggregate price index, such as the consumer price index (CPI), would be required to generate large fluctuations in output to offset the price-level effects from shocks to specific sectors. For instance, commodity-price movements tend to be both large and persistent, and influence the CPI directly through the price of gasoline and other forms of energy. Whereas a credible inflation-targeting central bank can generally look through these types of fluctuations, since their impact on inflation is highly transitory, a price-level-targeting central bank must respond by generating offsetting price-level movements in other sectors. As a result, price-level targeting could lead to greater aggregate volatility in an economy that is subject to large relative-price shocks.

This article reviews recent Bank of Canada research on the relative merits of price-level targeting (PLT) and inflation targeting (IT) for a small open economy that is subject to large and persistent terms-of-trade shocks.[2] The first section describes the basic mechanics of so-called history-dependent monetary policy, of which PLT is one special case, and discusses the

1 The potential benefits to the Canadian economy of reducing the inflation target from its current level of 2 per cent per year are also being explored.

2 A more general review of research on price-level targeting is provided in Ambler (2009).

conditions required for such policies to be beneficial in terms of economic stabilization. This is followed by an examination of recent model-based research comparing PLT with IT in a small open economy that is subject to relative-price shocks. The robustness of these results to, among other things, alternative assumptions about expectations formation is then discussed. Finally, research on the optimal price index to target under PLT is summarized before conclusions are drawn.

PLT as a Special Case of History-Dependent Monetary Policy

Targeting the price level, as opposed to the rate of inflation, can be thought of as a particular example of what is referred to as history-dependent monetary policy (Woodford 2003). History dependence simply means that monetary policy responds to past economic conditions, in addition to current and expected future conditions. This typically implies that policy will continue to respond to shocks, even after their impact on inflation and/or the output gap has fully dissipated. As a result, inflation will often exhibit a *secondary cycle*, meaning that the price-level effects generated by the shock will be partially, or fully, reversed.[3] For example, if a shock initially causes inflation to rise above some target rate, policy will continue to maintain interest rates above neutral until inflation moves below the target. This would imply that monetary policy causes inflation to undershoot the target when inflation is initially above target, and vice versa.

Based on this description, it is easy to see how price-level targeting represents a special case of history-dependent policy. Consider a central bank that chooses to target a constant price level through time. Following an economic shock that initially raises the price level (and creates inflation), the bank will subsequently engineer a period of deflation until the overall price level returns to the desired level. This type of response pattern is equivalent to responding to the sum of the current, and all previous, rates of inflation.[4] The appeal of PLT within the class of history-dependent policies is its transparency and relative ease of communication.

3 The term secondary does not mean that the cycle is of secondary importance, but that it comes after a first cycle.

4 In fact, the price level in any period is proportional to the product of all past gross inflation rates, and approximately equal to the sum of all past net inflation rates, where the gross inflation rate from period t to period $t+n$ is $\frac{P_{t+n}}{P_t}$ and the net inflation rate is $\frac{P_{t+n}}{P_t} - 1$.

Having established the mechanics of history dependence, we next turn to the fundamental question of how a central bank might benefit from adopting such an approach to setting monetary policy. It is not immediately obvious why a central bank seeking to stabilize inflation would want to cause secondary cycles in inflation, since this is clearly destabilizing to the economy, *other things being equal*. The key insight from the literature on history dependence is that such a policy will not leave other things equal. Specifically, if expectations of future inflation, which influence current inflation, correctly take account of the secondary cycle in inflation, they will exert a stabilizing effect on current inflation. Indeed, any policy that causes inflation to be lower (higher) in the future will also cause current inflation to be lower (higher) when expectations are forward looking. Intuitively, a firm that is considering a price change in the current period, knowing that this change will have to be reversed in the next period, will have less incentive to institute the change.

The appeal of PLT within the class of history-dependent policies is its transparency and relative ease of communication.

To better understand the mechanics of the expectations channel, consider the simplest form of the so-called New Keynesian Phillips curve (NKPC):

$$\pi_t = \beta\pi_{t+1|t} + \lambda y_t + \varepsilon_t, \qquad (1)$$

where π_t is the rate of price inflation,[5] $\pi_{t+1|t}$ is the rate of inflation expected to prevail in the next period (conditional on period-t information), y_t is the per cent difference between real GDP and potential GDP (i.e., the output gap), β and λ are constant parameters that are set to one for simplicity, and ε_t is a random shock, sometimes interpreted as a change to firms' desired markup of price over marginal cost. The New Keynesian model is based on two crucial assumptions: (i) firms change prices only periodically, meaning that prices generally remain fixed for more than one period, and (ii) firms form their expectations about the future in a rational way. Since it is known that the chosen price will likely remain in effect for

5 The inflation target is assumed to be zero.

multiple periods, account is taken of both current and expected future demand conditions, which implies that aggregate inflation is a forward-looking variable.

For the purpose of this discussion, we assume that inflation is determined according to equation (1) and that the instrument of monetary policy is the output gap. Thus, equation (1) also describes how policy influences inflation. Finally, for simplicity, we assume that the central bank cares equally about stabilizing inflation around its target and output around its potential. We can therefore describe the preferences of the central bank in terms of the following simple loss function:

$$L = \sigma_\pi^2 + \sigma_y^2, \tag{2}$$

where σ_π^2 and σ_y^2 are, respectively, the variance of inflation (relative to the target) and output (relative to potential output).

First, suppose that the central bank seeks to minimize equation (2) by responding only to current inflation. We can therefore write the central bank's reaction function as $y_t = \theta\pi_t$. Since we are assuming that $\lambda = 1$ and that ε_t is the only type of shock in the economy, we will obtain the result that $\theta = -1$. Now suppose that the economy is faced with a two-period shock in which $\varepsilon_1 = 1$, $\varepsilon_2 = 0.5$, and is zero thereafter. The optimal response of the output gap and inflation in each period is plotted in **Chart 1** (example 1), and, as our optimal rule implies, one is just the mirror image of the other, and total loss equals 0.91.

But suppose we relax the assumption that the central bank can respond only to current inflation and, instead, assume that it sets the same value of the output gap in each of the first two periods. In this scenario, the response is consistent with a reaction function of the form $y_t = \theta(\pi_t + \pi_{t-1})$. The optimal level for the output gap is -0.5 in both periods, which results in a total loss of 0.75 (example 2 in **Chart 1**). The reason behind this interesting result is quite simple: the output gap set in period 2 affects inflation in periods 1 and 2 when inflation expectations are forward looking, whereas the output gap set in period 1 affects inflation only in period 1. In this sense, the central bank obtains a better inflation/output trade-off by committing to generating a larger output gap in period 2 and a smaller output gap in period 1, relative to the first example. Of course, such a desirable outcome is possible only if inflation expectations explicitly take account of future demand conditions.

Chart 1: Benefits of history dependence

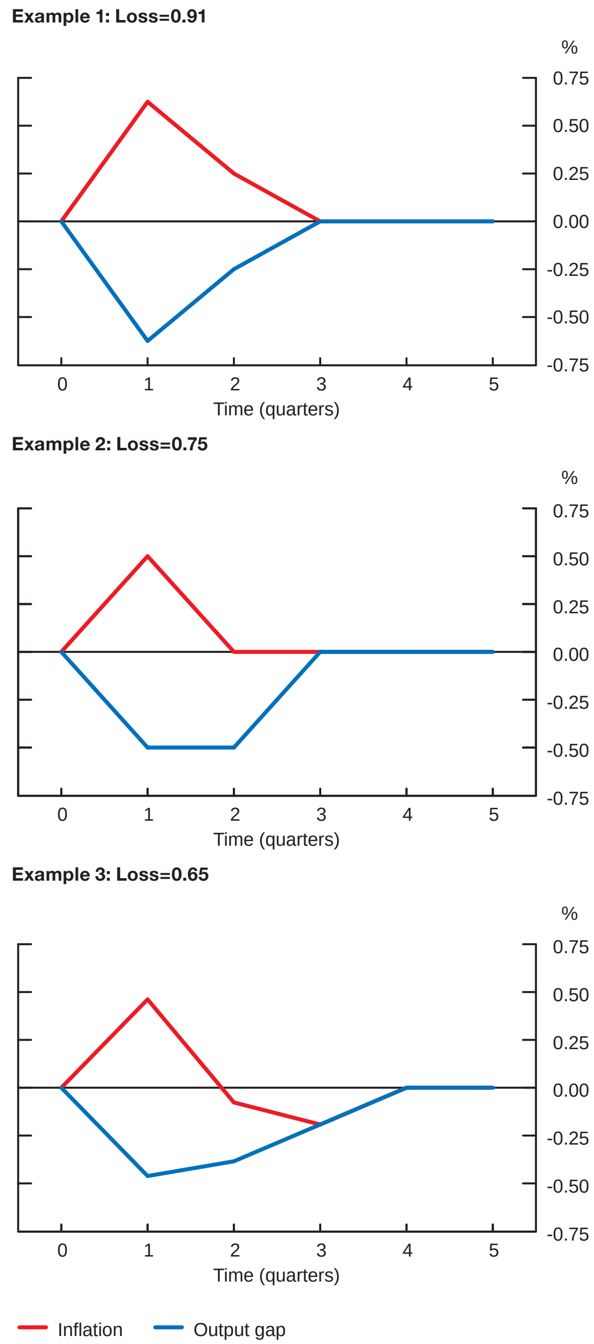

Source: Bank of Canada calculations

In this particular example, history dependence does not imply any undershooting of inflation, meaning that there are benefits to responding to past economic conditions even if no secondary cycle in inflation is generated. Nevertheless, an even better outcome can be obtained if a secondary cycle is permitted. For instance, suppose we now allow the central bank to choose the output gap as it wishes in each of the first 3 periods and that the output gap is zero thereafter (as shown in example 3, **Chart 1**). Given this option, the central bank generates a better inflation/output trade-off by maintaining the economy in excess supply in period 3, since this has a stabilizing effect on inflation in periods 1 and 2. The cost of this, as measured by deflation in period 3, is smaller than the benefit, since the overall loss declines from 0.75 in example 2, to 0.65 in example 3.

Woodford (2003) illustrates this basic point using the NKPC given by equation (1) and the loss function given by equation (2).[6] He shows that the optimal response to a positive markup shock, which initially causes inflation to rise, is to subsequently generate deflation until the price level returns to its pre-shock level. In other words, optimal monetary policy under commitment is consistent with targeting the price level, even though it is inflation that appears in the loss function. The particular policy rule consistent with achieving this outcome is given as[7]

$$y_t = y_{t-1} - \phi\pi_t, \tag{3}$$

which is history dependent in the sense that the central bank chooses the current period's output gap partly as a function of the previous period's output gap.[8]

That equation (3) implements PLT while setting the policy instrument as a function of inflation demonstrates the need to distinguish between policy regimes, such as IT and PLT, and the variables appearing in a history-dependent policy rule. In many instances, a history-dependent policy rule may implement aspects of both IT and PLT regimes in the short run. For instance, if we reduce the weight on the lagged output gap, y_{t-1} in equation (3) to a positive number less than one, then a positive markup shock may still eventually lead to a period of deflation, but it will be insufficient to fully return the price level to its control level. In this example, a deliberate undershooting of the inflation target may be inconsistent with the spirit of an IT regime, whereas not fully restoring the price level to its control level would be inconsistent with PLT.

As discussed in the next section, the grey area that exists between pure IT and pure PLT is quite important when researchers compare the two, using policy rules that feature interest rate smoothing.

Recent Research on Relative-Price Shocks and PLT

Comparisons of the efficacy of PLT relative to IT typically involve the use of optimized simple monetary policy rules that implement each regime in a quantitative macroeconomic model. This article surveys recent research using ToTEM, BoC-GEM, and a third small-open-economy model, all of which feature multiple production sectors and significant heterogeneity across sectors.[9]

The simple policy rules considered in each paper can be written as

$$R_t = \rho R_{t-1} + (1-\rho)R^* + \varphi_\pi(E_t\pi_{t+k} - \pi^T) + \varphi_y y_t,$$

(Inflation-forecast rule) (4)

for an IT regime, and

$$R_t = \rho R_{t-1} + (1-\rho)R^* + \varphi_p\big(E_t p_{t+k} - p^T_{t+k}\big) + \varphi_y y_t,$$

(Price-level-forecast rule) (5)

for a PLT regime, where R_t is the policy interest rate in period t; R^* is the long-run steady-state level of interest rates; $E_t\pi_{t+k}(E_t p_{t+k})$ is the period t expectation of inflation (log price level) in period $t+k$; y_t is the output gap; ρ, $\varphi_\pi(\varphi_p)$, and φ_y are fixed parameters that determine the degree of interest rate smoothing and the sensitivity of the policy rate to deviations of inflation (price level) from target and the output gap, respectively.[10] Note that the feedback horizon, k, determines the horizon of the response to inflation relative to its target, π^T, or the (log) price level relative to its target, p^T_{t+k}.

The first rule is referred to as an inflation-forecast (IF) rule, since the policy rate responds to a forecast of inflation, whereas the second rule is referred to as a price-level-forecast (PLF) rule. Since the IF rule ensures

6 Except that the weight on the variance of the output gap is less than one.

7 For simplicity, we ignore the initial-period problem in which policy does not respond to the lagged output gap. The issue of the time-inconsistency of this type of policy, as well as a suggested solution, is discussed in Woodford (2003).

8 If we solve equation (3) "backwards" to eliminate the lag of the output gap, we obtain a (negative) relationship between the current period's output gap and the sum of the current and all past inflation rates. This is equivalent to responding to the price level.

9 For a description of the Terms of Trade Economic Model, ToTEM, see Murchison and Rennison (2006). The Bank of Canada's version of the IMF's Global Economic Model, BoC-GEM, is described in Lalonde and Muir (2007), and a description of the third model can be found in de Resende, Dib, and Kichian (2010).

10 Inflation and interest rates are expressed as quarterly rates of change.

that the rate of inflation equals the target rate in the long run but, in general, makes no explicit provision to return the price level to a pre-specified level, this rule is loosely interpreted as implementing inflation targeting. The PLF rule, in contrast, does set policy explicitly to achieve a particular outcome for prices, $p_t = p_t^T$, and is therefore more consistent with price-level targeting in the long run. Having said that, just as the policy rule given by equation (3) implements aspects of both IT and PLT when the weight on the lagged instrument is less than one, the introduction of a lag of the instrument in equation (4) means that the IF rule will display history dependence and will therefore, to some degree, mimic the behaviour of a PLF rule with no lagged instrument. Similarly, equation (5) will, to some degree, mimic a rule that responds to the *sum* of past price-level gaps. As a result, some caution is warranted in mapping policy regimes, such as IT and PLT, to simple feedback rules such as the IF and PLF rules considered in these studies.

The version of ToTEM used in Murchison (forthcoming) explicitly models the CPI as a function of the Bank of Canada's measure of core CPI and the Canadian-dollar price of energy.[11] A permanent shock to the world oil price has both a demand component, driven by changes in wealth, and a relative-price channel, since commodities are both a factor of production of finished goods and final goods themselves (e.g., gasoline and home heating fuel). As a result, energy-price shocks involve a tension between stabilizing CPI inflation and stabilizing the output gap. Explicitly accounting for energy-price movements is crucial to the question addressed in Murchison since they explain much of the short-term volatility in the CPI, and their effect on the level of the CPI tends to be long lasting or permanent.

Murchison assumes that the policy-maker's preferences are well described by the following simple loss function:

$$L = \sigma_\pi^2 + \sigma_y^2 + 0.5\sigma_{\Delta R}^2, \qquad (6)$$

which penalizes the (unconditional) variance of CPI inflation and the output gap equally, and also puts a weight of 0.5 on the variance of the quarterly change in the policy interest rate, ΔR_t.[12] It is worth noting that this loss function accords no cost to price-level volatility per se, other than via its link to overall inflation volatility. Therefore, it does not capture any explicit benefits associated with reduced price-level uncertainty under PLT.

Using a distribution of shocks estimated by ToTEM over the period 1995Q1 to 2008Q4,[13] together with this loss function, the author simulates losses for different values of the policy-rule parameters ρ, $\varphi_\pi(\varphi_p)$, and φ_y for the IF and PLF rules described by equations (4) and (5). Those parameter values that produce the lowest value of L for each rule are retained and used to compare the IT and PLT regimes.

Coletti, Lalonde, and Muir (2008) use a very similar set-up but with a two-country (Canada and the United States), two-sector (tradables and non-tradables) version of the IMF's Global Economic Model (GEM), calibrated to Canadian and U.S. data from 1983 to 2004.[14] They also consider policy rules of the form given by equations (4) and (5) and a loss function similar to equation (6).

De Resende, Dib, and Kichian (2010) compare IT and PLT in an estimated small-open-economy model with multiple production sectors, sector-specific capital, and imperfect labour mobility between sectors. These model features are motivated by the idea that sector-specific shocks will generally mean that monetary policy will face a trade-off between stabilizing certain sectors and, consequently, destabilizing others. The importance of this trade-off will depend on the degree of factor mobility across sectors. While the authors also consider simple IF and PLF policy rules, their loss function is derived explicitly from the structure of the model.[15] As a result, the parameters of the policy rules are chosen to maximize the expected welfare of the representative household in the model, rather than an ad hoc loss function such as equation (6).

All three studies carefully consider the implications of relative-price shocks, including shocks that affect Canada's terms of trade, and broadly conclude that

11 The author uses the Bank of Canada's energy-commodity price index, which is converted to Canadian dollars using the current nominal exchange rate. This set-up assumes that movements in the world price of energy and the exchange rate are immediately and fully passed through to the consumer prices for energy products, such as gasoline, at a quarterly frequency.

12 Including ΔR_t in the loss function reduces the volatility of interest rate changes quite significantly but has little impact on the variance of inflation or the output gap. Excess instrument volatility may be disruptive to financial markets for reasons not captured by the models used.

13 The various types of structural shocks modelled in ToTEM are described in Murchison and Rennison (2006).

14 They estimate their shocks using a longer sample and use a smaller weight on the variance of the first difference of interest rates (0.1) than Murchison (0.5). They also use core CPI inflation in the loss function.

15 Welfare analysis is conducted based on a second-order approximation of the model (and the utility function) around its deterministic steady state.

Table 1: Inflation- and price-level-forecast rules

$$R_t = \rho R_{t-1} + \varphi_\pi E_t \pi_{t+k} + \varphi_p E_t p_{t+k} + \varphi_y \tilde{y}_t$$

		Coefficients of rule					Loss(PLF-IF)	Var. (PLF-IF)[a]		
Paper/Rule		ρ	φ_π	φ_p	φ_y	k	$\left(\frac{loss(PLF)}{loss\ (IF\ )} - 1\right)$	σ^2_π	σ^2_y	$\sigma^2_{\Delta R}$
Coletti, Lalonde, and Muir (2008)										
Unconstrained	IF	0.97	2.4	-	0.7	2	**-**	-	-	-
	PLF	0.85	-	3.7	0.9	3	**-1 %**	-2%	2%	0
de Resende, Dib, and Kichian (2010)[b]										
Unconstrained	IF	0.68	2.5	-	0.0	0	**-**	-	-	-
	PLF	0.0	-	1.1	0.0	0	**0%**	-	-	-
Constrained	IF	0.0	6.0	-	0.0	0	**-**	-	-	-
	PLF	0.0	-	1.1	0.0	0	**-5%**	-	-	-
Murchison (forthcoming)										
Unconstrained	IF	1.1	0.6	-	0.1	0	**-**	-	-	-
	PLF	0.98	-	0.09	0.2	4	**-5%**	-4%	-1%	0
Constrained	IF	0.0	3.75	-	0.3	1	**-**	-	-	-
	PLF	0.0	-	0.34	0.3	4	**-15%**	-7%	-1%	-7%
	IF	0.8	1.6	-	0.2	1	**-**	-	-	-
	PLF	0.8	-	0.1	0.2	4	**-9%**	-5%	-3%	-1%

a. Differences in variances across IF and PLF rules are expressed as a fraction of the total loss associated with the IF rule and weighted by their weight in equation (6). Thus, the differences for the three individual variables sum to the difference in loss (subject to rounding error).

b. Variances are not shown, since the differences in welfare-based loss cannot be expressed solely in terms of these variables.

PLF and IF rules yield very similar overall stabilization properties ("Unconstrained" rules, **Table 1**). When all types of shocks are considered, the PLF rule dominates the IF rule by a small margin in ToTEM and in BoC-GEM, whereas de Resende, Dib, and Kichian find no difference. In addition, the results for ToTEM and BoC-GEM suggest that when inflation expectations are calibrated to be highly forward looking, the PLF rule also dominates the IF rule in the presence of relative-price shocks.[16] In other words, the gain realized via the expectations channel outweighs the loss associated with having to stabilize the overall price level in response to sector-specific shocks.

The impact of a permanent 20 per cent increase in the world price of energy, simulated using ToTEM, is illustrated in **Chart 2**. Three policy rules are used: the optimized IF rule, the optimized PLF rule, and fully optimal policy under commitment (labelled Optimal). Fully optimal policy is a natural benchmark: it represents the absolute best outcome that policy can achieve for a given model and loss function.[17] Regardless of the rule considered, an unanticipated rise in energy prices causes an immediate increase in the Canadian-dollar price of energy and, hence, in the overall CPI **(Chart 3)**.

The transmission of commodity-price shocks in ToTEM is discussed extensively in Murchison and Rennison (2006). For the purposes of this article, it is sufficient to highlight that slightly more than 25 per cent of the increase in the world energy price is offset (with the IF rule), in terms of the Canadian-dollar price, by an immediate and permanent appreciation of the Canadian dollar. As a result, the overall increase in the CPI is more muted than would be the case with a fixed exchange rate. As the exchange rate appreciation is gradually passed through to import and export prices, net exports weaken, and upward pressures on core CPI inflation decline.

The responses of the IF and PLF rules to the shock are broadly similar: policy gradually tightens (years 1 and 2) and then loosens, in both cases by a modest

16 Coletti, Lalonde, and Muir (2008) also consider a more recent sample (1995 to 2006), nearly identical to that used by Murchison, in which the persistence of inflation is lower than over their full sample. As a result, the weight on lagged inflation in their NKPC is set to zero, and PLT dominates IT for all shocks, including relative-price shocks.

17 Optimal policy is computed in ToTEM following Dennis (2007). It is optimal under the assumption of no uncertainty other than that associated with imperfect knowledge of future shocks and the loss function given by equation (6). It would not generally be optimal in the presence of non-additive uncertainty, such as parameter, model, and real-time data uncertainty (Cateau and Murchison 2010).

Chart 2: Results of a permanent 20 per cent increase to the world price of energy in ToTEM

a. Total inflation

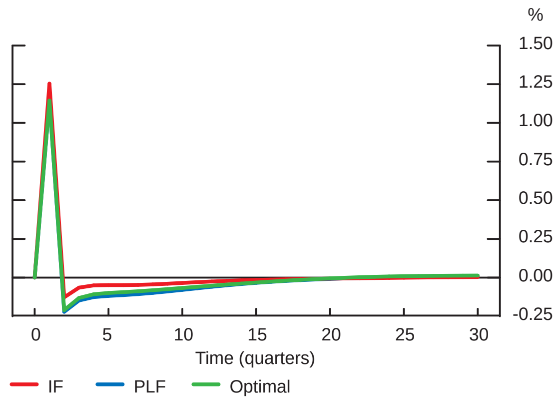

b. Output gap

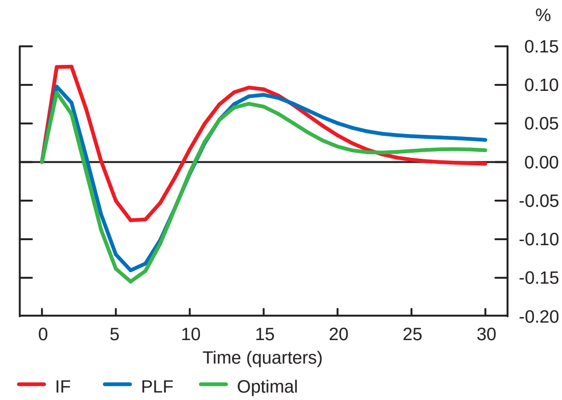

c. Core CPI inflation

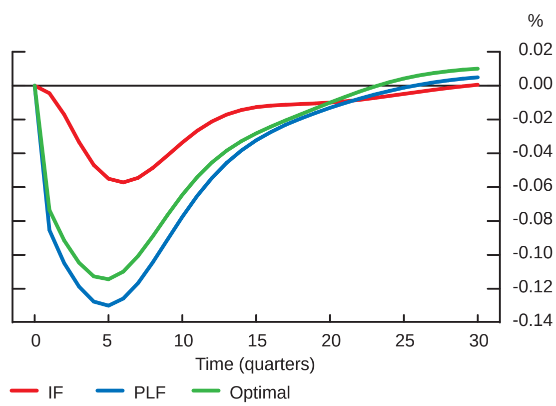

d. Log CPI level

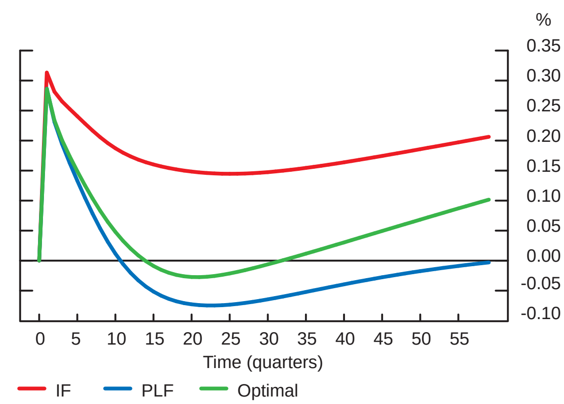

e. Policy rate

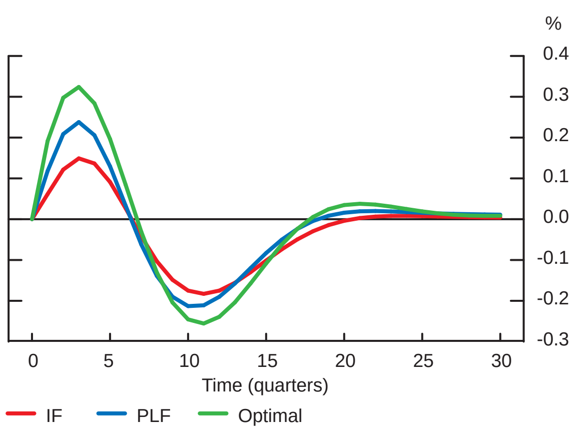

f. Real exchange rate

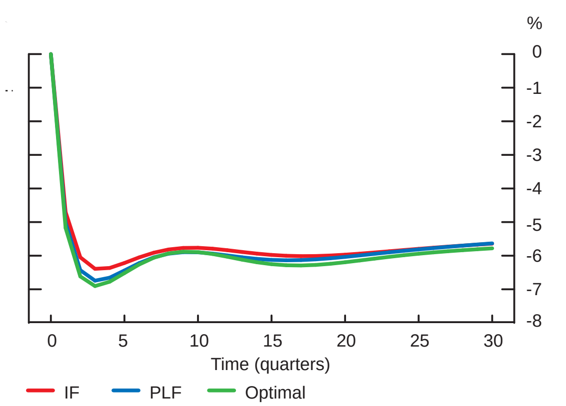

Note: A decrease represents an appreciation in the real exchange rate.

Source: Bank of Canada calculations

Chart 3: Energy prices

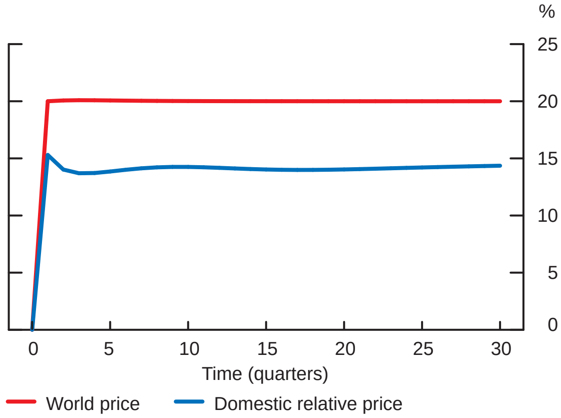

Source: Bank of Canada calculations

amount.[18] However, under the PLF rule, it tightens by roughly 50 per cent more at the peak, implying a smaller initial increase in the output gap and a sharper decline in core CPI inflation.[19] Higher real interest rates also cause a more pronounced appreciation of the exchange rate with the PLF rule, meaning that a smaller proportion of the increase in world energy prices gets passed on to Canadian consumers, and that net exports are weaker in year 2, relative to the IF rule.

Based on these simulation results, the intuition that PLT requires greater volatility in output to stabilize the price level in response to a terms-of-trade shock is validated by ToTEM. To restore the CPI to the target, the PLF rule creates roughly twice as much excess supply (at the trough) as the IF rule. In terms of overall loss, however, which also factors in CPI inflation and instrument volatility, the PLF rule (PLT) still outperforms the IF rule (IT) by 4 per cent, because the initial rise in inflation is smaller under PLT.[20]

To summarize, simulations with ToTEM find that PLT is well suited to handling energy-price shocks and relative-price shocks more generally[21] and that it comes very close to replicating fully optimal policy under commitment. In fact, **Chart 2** illustrates that while optimal policy represents a hybrid between the IF and PLF rules, in the short run, it follows the PLF rule much more closely. It is not until year 3 that optimal policy takes a more expansionary stance, thereby allowing the overall price level to rise permanently above control. Under the PLF rule, the CPI returns to the target near the end of year 3 but then remains below the target for several years. This undershooting of the price level is due to the high weight $(\rho = 0.98)$ on the lagged interest rate in the PLF rule and represents another example of the effect of introducing history dependence.

Simulations with ToTEM find that PLT is well suited to handling energy-price shocks and relative-price shocks more generally and that it comes very close to replicating fully optimal policy under commitment.

When all types of shocks are considered, Murchison shows that the median time required for the price level to return to target is substantially longer than the target horizon for inflation under IT, when optimized simple policy rules are considered. Specifically, it is shown that in a stochastic environment, with representative shocks drawn from the 1995–2008 sample hitting the economy each quarter, the median time required to return the price level to within ±0.5 per cent of the target is about 2.5 years,[22] about double that required to return year-over-year inflation to within ±0.1 percentage points of the target with an optimized IF rule.

As discussed in the previous section, responding to past economic conditions implies history dependence, which can have an important stabilizing effect on the economy when expectations explicitly take into account this feature of monetary policy. History dependence can be introduced directly, via the inclusion of lagged inflation in the policy rule (see example 2, **Chart 1**), or by responding to lags of the policy instrument itself (as in equations 3, 4, and 5). In all three studies cited here, the optimized IF rules respond positively to the level of the policy interest rate in the previous quarter, and the weights (captured by the parameter ρ) range from 0.68 to 1.1. In other words,

18 The shock pushes the Canadian economy into modest excess demand for about one year after the shock. As a result, even the IF rule initially tightens policy, despite the decline in pressures on core CPI inflation.

19 For simplicity, the price-level target in the PLF rule and the inflation target in the IF rule are both zero.

20 While the difference in the initial rise in inflation between the IF and PLF rules is small, loss is calculated using the squared deviation of inflation for each rule. Therefore, the larger the overall inflation response, the greater will be the loss difference, for a given difference in responses across rules. In this shock, the CPI inflation responses peak at between 1.0 and 1.25 percentage points above control, expressed at annual rates.

21 See Murchison (forthcoming) for additional examples.

22 Under the assumption of no future shocks. The choice of 0.5 per cent as the threshold is arbitrary but seems reasonable considering the unconditional variance of the price level under PLT, using the optimized PLF rule.

the similarity between the performance of the IF and PLF rules found in these studies may be partly due to the fact that an IF rule with a high value of ρ can closely mimic the behaviour of a PLF rule.

To explore the sensitivity of the results to the degree of interest rate smoothing, Murchison (forthcoming) and de Resende, Dib, and Kichian (2010) also compare optimized IF and PLF rules that restrict ρ to zero, while Murchison also explores rules with ρ equal to 0.8, which corresponds to the average of the historical estimates for Canada ("Constrained" in **Table 1**). In all cases, eliminating history dependence via interest rate smoothing penalizes the IF rules somewhat more than the PLF rules. Put a different way, rules that are already history dependent, owing to the inclusion of the price level, benefit relatively less from the additional history dependence introduced through the lagged interest rate term.

De Resende, Dib, and Kichian report that their preferred PLF rule generates a 5 per cent reduction in loss relative to the IF rule when $\rho = 0$, compared with no difference when $\rho > 0$. The corresponding numbers for Murchison are somewhat larger—15 per cent ($\rho = 0$) and 9 per cent ($\rho = 0.8$)—since interest rate volatility is explicitly penalized in equation (6) but does not generally appear in welfare-based loss functions.[23]

To summarize: When monetary policy commits to setting the current policy rate partly as a function of the past rate, in addition to the output gap and to a forecast of either inflation or the price level, then IF and PLF rules are fairly similar in terms of their economic-stabilization properties. When policy is restricted to responding only to the output gap and a forecast of either inflation or the price level, then PLF rules are found to dominate IF rules. This suggests that there may be modest economic gains, measured in terms of greater stability, associated with the adoption of a target for the price level rather than for the rate of inflation.

Other Considerations

Robustness

The discussion so far has emphasized the crucial link between the performance of history-dependent monetary policy, including PLT, and the presence of forward-looking price-setting behaviour in the economy. Steinsson (2003) shows that as the relative importance of forward-looking expectations declines in the economy, so do the benefits of fully returning the price level to control following a markup shock. Coletti, Lalonde, and Muir (2008) confirm the same basic result, using a more realistic quantitative model: the relative performance of an optimized PLF rule depends importantly on the weight on lagged inflation. This result is quite intuitive: when pricing decisions depend on past, as opposed to future, economic conditions, future monetary policy actions become less influential for current price-setting behaviour.

In a follow-up paper, using a version of BoC-GEM that explicitly models emerging Asia and the block of commodity-exporting countries, Coletti et al. (forthcoming) show that when inflation is partially backward looking and the short-run supply and demand curves for energy are highly inelastic, IT dominates PLT in response to energy-price shocks, albeit by a modest amount. They also explore the idea that the source of the shock driving the terms of trade may matter for comparisons of PLT and IT. For example, the authors also consider the impact of a permanent increase in global productivity on commodity-importing regions. This shock has important implications for both the price of Canada's exports (through higher energy prices) and for the price of imported goods (through a stronger exchange rate). In this instance, IT outperforms PLT by a significant margin, close to 25 per cent, which is substantially larger than in the case of an oil-supply shock. This is explained by two factors. First, in this version of BoC-GEM, a permanent shock to the demand for oil induces a more persistent response in the price of oil and in marginal cost than a permanent shock to the supply of oil. Second, as opposed to a supply shock, a demand shock increases both the price of oil and the price of non-energy commodities, which reinforces the effect of the shock on the marginal cost. As a result, the impact on marginal cost is larger and more persistent for a demand shock than a shock to the supply of oil. Given the very different results across the different types of shocks to the terms of trade, it would be very useful to have a better idea of the relative importance of these types of shocks for the Canadian economy.

Murchison (forthcoming) generalizes these results somewhat, showing that as past economic conditions become relatively more important than future conditions to current private sector decisions, the relative performance of PLT tends to diminish, since the expectations channel becomes relatively less influential.[24]

23 Responding to the lagged interest rate introduces additional inertia in interest rates, which reduces the variance of interest rate changes.

24 Short-run adjustment costs, rule-of-thumb behaviour, and habit persistence in consumption all tend to increase the relative importance of past economic conditions.

For instance, when households place a high weight on smoothing the growth rate of consumption, the level of previous consumption becomes a more important determinant of current consumption, and the future path of real interest rates becomes relatively less important. Similarly, as short-run adjustment costs associated with changing the relative intensities of factor inputs, such as installed capital, increase, the level of the capital stock in the previous period becomes a more important determinant of the current capital stock.

The overall robustness of PLT will depend on all of the structural parameters that govern the dynamics of the model in question, as well as the overall degree of uncertainty regarding their true values. In a related paper, Cateau, Desgagnés, and Murchison (forthcoming) derive optimized inflation- and price-level-forecast rules for ToTEM and compare their performance across 5000 different parameterizations of the model.[25] They conclude that, overall, optimized PLF rules are more robust to this form of uncertainty than optimized IF rules.

What is the appropriate price index to target?

In a simple one-good model with no relative prices, the choice of the price index is trivial. However, in more realistic multi-good models, such as those reviewed here, the question of what constitutes an ideal price index to target in a PLT regime can be considered from the perspective of minimizing either an ad hoc loss function, such as equation (6), or a welfare-based loss function. De Resende, Dib, and Kichian (2010) compare the performance of simple rules across five distinct sectoral price indexes—the consumption sector (CPI), non-tradables, tradables, manufacturing, and import prices—and find that targeting the CPI maximizes household welfare. Indeed, CPI targeting comes quite close to replicating the level of welfare that would obtain in the absence of nominal-price rigidity. The authors attribute this result to the inclusion of capital-adjustment costs in their model. Specifically, they show that when the cost of adjusting the capital stock in the non-tradable goods sector is low, it is optimal to target the price level in this sector. This result is consistent with previous work in the literature (Erceg, Henderson, and Levin 2000), which shows that monetary policy should aim to stabilize the price level in the sector with the stickiest prices, since it is precisely this stickiness that leads to suboptimal resource allocation and, hence, reduced welfare.[26] De Resende, Dib, and Kichian show that this result need not hold when other sources of rigidity are included in the model.

Shukayev and Ueberfeldt (2010) go a step further and compute the index weights for the eight major sub-components of the CPI that maximize the expected utility of the representative household in their model. In theory, these weights could differ substantially from the expenditure-based weights used by Statistics Canada if there are significant differences in price stickiness across the various components of the CPI. Using a model that includes sector-specific shocks to productivity and price markups, they find the welfare gain from using a PLF rule that responds to the ideal index, relative to the expenditure-based index, to be small.

Conclusion

This article reviews recent Bank of Canada research on the relative merits of price-level targeting and inflation targeting for a small open economy that is subject to large and persistent terms-of-trade shocks. While the quantitative results are mixed and somewhat dependent on the specific features of the model employed and the calibration of expectations, the balance of evidence suggests that PLT and IT, implemented through simple PLF and IF rules, are fairly similar in their ability to stabilize inflation, the output gap, and interest rates, although PLF rules generally perform better. Furthermore, this conclusion is robust to the inclusion of several types of relative-price shocks, including shocks to the terms of trade, although the results in Coletti et al. (forthcoming) indicate that the underlying source of terms-of-trade movements may matter for this assessment. Finally, the research suggests that, conditional on adopting PLT, the overall CPI would represent close to an ideal index to target.

25 These parameters are drawn from the Bayesian posterior distribution of the estimated parameters.

26 The basic intuition for this stylized result is straightforward: if monetary policy can fully stabilize the price level in that sector, the welfare consequences of nominal rigidity become zero, because firms have no incentive to change prices.

Literature Cited

Ambler, S. 2009. "Price-Level Targeting and Stabilization Policy: A Review." *Bank of Canada Review* (Spring): 19–29.

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

Cateau, G., H. Desgagnés, and S. Murchison. Forthcoming. "Monetary Policy Rules in an Uncertain Environment." Bank of Canada.

Cateau, G. and S. Murchison. 2010. "Monetary Policy Rules in an Uncertain Environment." *Bank of Canada Review* (Spring): 27–39.

Coletti, D., R. Lalonde, and D. Muir. 2008. "Inflation Targeting and Price-Level-Path Targeting in the Global Economy Model: Some Open Economy Considerations." *IMF Staff Papers* 55 (2): 326–38.

Coletti, D., R. Lalonde, P. Masson, D. Muir, and S. Snudden. Forthcoming. "Commodities and Monetary Policy: Implications for Inflation and Price-Level Targeting." Bank of Canada.

Dennis, R. 2007. "Optimal Policy in Rational Expectations Models: New Solution Algorithms." *Macroeconomic Dynamics* 11 (01): 31–55.

de Resende, C., A. Dib, and M. Kichian. 2010. "Alternative Optimized Monetary Policy Rules in Multi-Sector Small Open Economies: The Role of Real Rigidities." Bank of Canada Working Paper No. 2010-9.

Erceg, C., D. Henderson, and A. Levin. 2000. "Optimal Monetary Policy with Staggered Wage and Price Contracts." *Journal of Monetary Economics* 46 (2): 281–313.

Lalonde, R. and D. Muir. 2007. "The Bank of Canada's Version of the Global Economy Model (BoC-GEM)." Bank of Canada Technical Report No. 98.

Murchison, S. Forthcoming. "Consumer Price Index Targeting." Bank of Canada.

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

Shukayev, M. and A. Ueberfeldt. 2010. "Price Level Targeting: What Is the Right Price?" Bank of Canada Working Paper No. 2010-8.

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

Woodford, M. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton, NJ: Princeton University Press.