

Bank of Canada



Banque du Canada

Working Paper 2001-24 / Document de travail 2001-24



Price-Level versus Inflation Targeting in a Small Open Economy

by

Gabriel Srouf

ISSN 1192-5434

Printed in Canada on recycled paper

Bank of Canada Working Paper 2001-24

December 2001

Price-Level versus Inflation Targeting in a Small Open Economy

by

Gabriel Srouf

Research Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
gsrouf@bankofcanada.ca

The views expressed in this paper are those of the author.
No responsibility for them should be attributed to the Bank of Canada.

Contents

Acknowledgements.....	iv
Abstract/Résumé.....	v
1. Introduction.....	1
2. The Closed-Economy Case.....	3
2.1 The baseline model.....	3
2.2 The basic argument.....	5
2.3 Output variability.....	8
2.4 The role of forward-looking behaviour.....	12
2.5 The infinite-horizon model.....	13
3. Small Open Economy.....	15
4. Conclusion.....	18
References.....	20

Acknowledgements

The author would like to thank Allan Crawford, Brian O'Reilly, Nicholas Rowe, Larry Schembri, and seminar participants at the Bank of Canada for helpful comments. Of course, the author is solely responsible for any remaining errors.

Abstract

This paper compares two types of monetary policy: price-level targeting and inflation targeting. It reviews recent arguments that favour price-level targeting, and examines how certain factors, such as the nature of the shocks affecting the economy and the degree to which agents are forward-looking, bear upon the arguments. The paper then extends the analysis to a small open economy such as Canada's, and considers whether it is practical for this country to pursue price-level targets if its dominant trading partner, the United States, allows the price level to drift.

JEL classification: E52

Bank classification: Monetary policy framework

Résumé

L'auteur compare deux types de politique monétaire, l'un fondé sur la poursuite de cibles définies à l'égard du niveau des prix et l'autre sur celle de cibles d'inflation. Faisant la revue des arguments invoqués récemment en faveur des cibles de niveau des prix, il examine comment certains facteurs, tels que la nature des chocs qui secouent l'économie et la mesure dans laquelle les agents sont tournés vers l'avenir, influent sur la validité de ces arguments. Puis il étend l'analyse au cas d'une petite économie ouverte comme le Canada et se demande si ce pays peut envisager, de façon réaliste, d'adopter des cibles formulées en fonction du niveau des prix alors que son partenaire commercial dominant, les États-Unis, tolère une dérive du niveau des prix.

Classification JEL : E52

Classification de la Banque : Cadre de la politique monétaire

1. Introduction

There is a wide consensus today that a primary objective of monetary policy ought to be price stability. However, the debate continues as to exactly what price stability should mean: a low and stable *inflation rate*, or a stable *price level*? For now, most countries have opted for a low and stable inflation rate. Canada, in particular, has chosen to target the inflation rate since 1991. But, as a low and stable inflation rate has successfully been secured in many countries, and more experience has been acquired in this environment, there have been some calls for monetary authorities to take an extra step and target the price level.

The conventional view is that a policy that targets the price level generates too much volatility in output and prices in the short run, because it requires that prices be returned to their previous level after every shock. Of course, in the long run, prices are fully predetermined. On the other hand, a policy that targets the inflation rate is thought to produce less volatility in the short run, because it accommodates unanticipated shifts in prices, and requires only that the inflation rate be returned to its previous level. As a consequence, however, prices are highly uncertain in the long run. Thus, the choice between price-level and inflation targeting in the past was framed as a choice between predetermined prices in the long run and stability in the short run.

Recently, a number of authors have disputed this view.¹ They contend that, contrary to conventional thinking, targeting the price level induces low volatility in prices, both in the long run and in the short run. Their rationale is that, under price-level targeting, agents would expect prices to always return to their initial level, and therefore would be reluctant to change their own prices following shocks. In other words, price-level targeting would inhibit price changes from occurring in the first place. Accordingly, these authors support a certain degree of price-level stability in monetary policy.

The purpose of this paper is twofold. First, it reviews the argument in favour of price-level targeting in some detail. Second, since the literature has focused on closed economies, this paper examines whether a similar argument applies to a small open economy such as Canada's. In particular, the paper inquires whether it is practical for Canada to pursue price-level targets if its dominant trading partner, the United States, does not.

The paper is organized as follows. Section 2 reviews the argument in favour of price-level targeting in a simple closed-economy model. Following common practice, overall welfare in the

1. See, for instance, Duguay (1994), Svensson (1999), Svensson and Woodford (1999), Vestin (1999), and Woodford (1999).

economy is defined solely in terms of deviations of inflation and output from equilibrium:² the smaller the deviations, the larger the welfare. The first-best policy is then derived in this context, and compared alternatively to inflation targeting and price-level targeting. The argument in favour of price-level targeting is most convincing when expectations have a strong influence on current prices, and deviations in output needed to move prices back to their initial level following shocks are not costly. Thus, it is first shown that if the Phillips curve is purely forward-looking, and output variability does not affect welfare, then the first-best policy in fact coincides with price-level targeting.

In general, however, the first-best policy involves only a partial return of prices to their initial level following shocks. On one hand, the first-best policy always involves some movement of prices back towards their initial level (unless the Phillips curve is purely backward-looking), because, at the margins, the welfare cost of the necessary output deviation is of second order relative to the welfare gain of a movement back in prices.³ On the other hand (unless the Phillips curve is purely forward-looking), the first-best policy does not fully return prices to their initial level. The intuitive reason is that, if current inflation does not depend one-for-one on expected future inflation, then a full return of prices to their initial level following a shock would reduce the immediate effect of the shock on contemporaneous inflation by an amount smaller than the subsequent deflation induced by the movement back in prices. The trade-off would therefore not be worthwhile. The extent and the rate at which the first-best policy moves prices back towards their initial level following shocks, and hence the relative merits of price-level targeting, depend on all the model's specifications (including the data-generating process of all the variables), and are highly empirical questions.

Section 3 extends the analysis to the case of a small open economy with a flexible exchange rate. We show first in a very simple model that the same conclusions reached in the case of a closed economy apply in a small open economy. We then argue that the analysis may be too simplistic for a country such as Canada, and that it is subject to important qualifications. In particular, the exchange rate is likely to be highly volatile if Canada and the United States pursue very different monetary policies, and such volatility could be harmful to the economy. Section 4 concludes.

This paper follows a string of other papers on the same subject.⁴ One unique feature of this paper is that the basic argument is set in a context broad enough to subsume various special cases, and

-
2. More precisely, the social-loss function is defined as a weighted sum of squared deviations of inflation and output from equilibrium.
 3. This is a consequence of the quadratic form of the loss function.
 4. In addition to the authors cited in footnote 1, see Clarida, Gali, and Gertler (1999), Dittmar and Gavin (2000), and Barnett and Engineer (2000). There is, in particular, a strong overlap between the latter paper and ours, although, except for section 2.5, our paper was written independently. Section 2.5 was added for completeness after reading Barnett and Engineer. A first version of our paper was available in May 2000.

simple interpretations of the results are provided. To our knowledge, the open-economy case has not yet been examined in the literature.

2. The Closed-Economy Case

The argument that confers comparative advantage to price-level targeting is clearest when the economy is confined to two periods: prices are subject to shocks in period 1 and are influenced by expected prices in period 2. Accordingly, this section considers first a two-period-horizon model of a closed economy. It derives the first-best policy and compares it with price-level targeting and inflation targeting. Because, through expectations, price-level targeting can affect not only prices but other state variables as well, abstraction is first made of the effects of expectations on other state variables. More general conditions, including the case of an infinite horizon, are discussed subsequently.

In this section, particular care is taken to ensure that the results will extend easily to the case of a small open economy examined in section 3. In a sense, this section does most of the groundwork for section 3.

2.1 The baseline model

Suppose that prices in the economy are governed by a standard Phillips curve of the form

$$\pi_t = \beta\pi_{t+1|t} + BZ_{t|t-1} + \varepsilon_t, \quad (1)$$

where β is a discount rate, π_t is the inflation rate, Z_t is a vector of variables, and ε_t is a white-noise shock that is observed after monetary action at time t is taken.⁵

Although the notation is slightly tedious, this formulation is very useful, for it subsumes a variety of specifications, and will be most convenient in the sequel to extend the results beyond the base case. For example, the model can allow for lagged variables, which can be incorporated in the vector Z_t , or even contemporaneous variables, which can always be written as the sum of an anticipated term (that can be incorporated in the vector Z_t) and an unanticipated term (that can be incorporated in the shock ε_t).

The value of $BZ_{t|t-1}$ is assumed to be determined by the expected value, $i_{t|t-1}$, of the monetary instrument at time t . Therefore, the value of $BZ_{t|t-1}$ is essentially controlled by the central bank. This applies, for instance, if Z_t consists solely of the output gap, y_t , and $y_{t|t-1}$ is determined by

5. As usual, $x_{i|j}$ denotes the expected value of the variable x at time i based on information at time j .

$i_{t|t-1}$.⁶ The expectational form of $BZ_{t|t-1}$ conforms with the conventional wisdom that prices are sticky and therefore respond with a lag to unanticipated shocks.

The crucial element in the present model is the forward-looking expectation, $\pi_{t+1|t}$. It expresses the fact that expected future changes in prices feed into current pricing decisions, and implies that future monetary policy can affect current prices. For example, according to this model, expectations of a contractionary monetary policy, and hence lower inflation in the future, induce lower inflation in the current period.

For simplicity, the model is assumed to effectively range over only two periods. Specifically, the economy is assumed at time $t = 0$ to be in steady state, whereby the variables π_t and Z_t , and their expected future values, equal 0.⁷ The economy is then subject to a single unanticipated inflationary shock, ε_1 , at time 1. At time 3, the economy automatically returns to the steady state under all circumstances, and monetary policy has no further bearing on the economy. (The infinite-horizon model is examined in section 2.5.)

Consistent with the specification of a two-period horizon, the social welfare function of the economy at time t is assumed to have the form

$$-L_t = -E_t[(\pi_t^2 + \lambda y_t^2) + \beta(\pi_{t+1}^2 + \lambda y_{t+1}^2)], \quad (2)$$

where y_t is the output gap, and λ is a weight on output variability, $0 \leq \lambda \leq \infty$.⁸ According to this function, only inflation and output variability are detrimental to the economy, and price-level variability per se is not.

2.1.1 Terminology

The policy that maximizes the welfare function at time 1, when the central bank can make precommitments about its future actions, is called the *first-best policy*. In general, this policy is not time-consistent. The time-consistent policy is called simply *inflation targeting*.⁹ That is,

-
6. More generally, the assumption applies if $BZ_{t|t-1} = bi_{t|t-1} + CX_{t|t-1}$, ($b \neq 0$), and the central bank takes action at time t after it observes all variables in $X_{t|t-1}$.
 7. The price level at time 0 is also assumed to equal 0.
 8. The inflation target is assumed to equal 0. However, the welfare function can be trivially modified to accommodate a positive inflation target.
 9. Sometimes in the literature, the first-best policy is called inflation targeting with precommitment, and the time-consistent policy is called inflation targeting without precommitment. We prefer to reserve the term “inflation targeting” solely for the time-consistent policy, because we believe that is what is usually meant by inflation targeting in practice.

inflation targeting is the policy that applies when the decision-maker seeks to maximize the welfare function (2) *in every period*.

Likewise, the time-consistent policy that seeks, *in every period*, to maximize the function

$$-E_t[(p_t^2 + \lambda y_t^2) + \beta(p_{t+1}^2 + \lambda y_{t+1}^2)], \quad (3)$$

where p_t is the price level, is called simply *price-level targeting* (rather than price-level targeting without precommitment).

Section 2.2 shows that, under some circumstances, the first-best policy is akin to price-level targeting, and that, consequently, a central bank can adopt price-level targeting as a proxy for the first-best policy. This begs the question, cannot a central bank just adopt the first-best policy? The answer emphasized in the literature refers to the rule-versus-discretion type of problem, which argues that the monetary authorities cannot credibly make the precommitment necessary to achieve the first-best policy, because they have an incentive to maximize the welfare function anew every period, and thus to ignore past commitments. The pursuit of price-level targeting, which by definition is time-consistent, would then provide a way to circumvent this problem. This proposition is analogous to Rogoff's finding that the appointment of a conservative central banker can provide a way to circumvent an inflation bias à la Barro-Gordon (Rogoff 1985).

The rationale just described has met two criticisms. One is that a central bank would not renege on its past commitments, knowing that this would ultimately harm welfare (McCallum 1997). The other is that, if a central bank cannot make credible commitments, then it should not be able to commit to price-level targeting.

We offer an alternative answer, which relies on the importance of communication and transparency in public policy. We submit that time-consistency is highly desirable in public policy, for it means that the conduct of policy and its objective can be explained clearly and simply to the public at any time. Time-*in*consistency can make a policy impractical, for it requires that monetary authorities explain their future conduct in terms of their past behaviour. Thus, price-level targeting, which is time-consistent (relative to the objective function (3)), can provide a practical way to implement the first-best policy, which is time-*in*consistent.

2.2 The basic argument

As the social welfare function (2) indicates, the smaller the deviations of inflation and output from equilibrium over periods 1 and 2, the greater the social welfare. According to the Phillips curve (1),

inflation in period 1, π_1 , is the result of the combined effect of the contemporaneous shock, ε_1 , and expected inflation in period 2, $\pi_{2|1}$. Inflation in period 1 would therefore react less to the shock ε_1 , and welfare would be enhanced, if prices were expected to revert to their initial level in the next period.

This argument is most convincing when price stability is the only concern (i.e., $\lambda = 0$), since in that case moving prices back to their initial level in period 2 does not carry any cost in terms of output. Indeed, we will show that the first-best policy coincides with price-level targeting when $\lambda = 0$.

In general, moving prices back towards their initial level in period 2 entails some welfare cost in terms of output variations. At the margins, that cost is always smaller than the welfare gain that results from reducing inflation in period 1. Therefore, under very general conditions, the first-best policy will always involve some movement of prices back towards their initial level. However, the extent and the rate at which this occurs will depend on the model's specifications, including the weight placed on output variability in the welfare function, and the data-generating process governing output. This proposition is examined in more formal terms below.

We first consider the case where output variability does not affect economic welfare (i.e., $\lambda = 0$), and the social loss function has the form:

$$L_t = E_t[\pi_t^2 + \beta\pi_{t+1}^2]. \quad (4)$$

The discount rate β actually cancels out in all the derivations (with the two-period model). Henceforth, unless stated otherwise, β is assumed to equal 1.

Note first that, since the economy is assumed to be in steady state at time 0 and at time 3, $BZ_{1|0} = 0$ and $\pi_{3|2} = 0$. Also, $\varepsilon_2 = 0$, since there are no shocks after time 1. From the Phillips curve (1) it follows therefore that

$$\pi_1 = \pi_{2|1} + \varepsilon_1 \quad (5)$$

and

$$\pi_2 = BZ_{2|1}. \quad (6)$$

2.2.1 *The first-best policy*

Since π_1 depends one-for-one on $\pi_{2|1}$ (equation (5)), the social cost at time 1, $L_1 = E_1[\pi_1^2 + \pi_2^2]$, is minimized when the inflation rate is smoothed equally between the two periods. That is,

$$\pi_{2|1}^2 = \pi_1^2 \quad (7)$$

which is to say that

$$\pi_{2|1} = -\pi_1. \quad (8)$$

This applies if $BZ_{2|1}$ is set equal to $-\pi_1$. Substituting (8) back into equation (5) gives

$$\pi_1 = \frac{\varepsilon_1}{2}. \quad (9)$$

Thus, under the first-best policy, only half of the shock is assimilated into prices in period 1, although at the cost of a change in prices of equal magnitude, but opposite sign, in period 2. The overall loss at time 1 amounts therefore to

$$L_1 = \frac{1}{2}\sigma_\varepsilon^2, \quad (10)$$

where σ_ε is the standard deviation of the random shock, ε .

2.2.2 *Inflation targeting*

In contrast, under inflation targeting, prices would not be expected to fall in period 2, since the rational public would expect the authorities to target 0 inflation at that time. Hence,

$$\pi_{2|1} = 0, \quad (11)$$

which applies if $BZ_{2|1}$ is set equal to 0.

In that case, the full shock is assimilated into prices in period 1,

$$\pi_1 = \varepsilon_1, \quad (12)$$

and the overall loss at time 1 is twice as large as that under the first-best policy:

$$L_1 = \sigma_\varepsilon^2. \quad (13)$$

2.2.3 Price-level targeting

The first-best outcome also applies if the authorities are believed to target the price level instead of the inflation rate; i.e., if they are believed to maximize in every period equation (3) instead of equation (2) (again with $\lambda = 0$). Indeed, under this policy, the objective at time 2 is precisely to reverse the shock to inflation at time 1, to bring the price level back to its initial level.

We can make the following conclusion:

In the two-period model, when output variability is not a concern for welfare, the first-best policy coincides with price-level targeting. (We shall see in section 2.5 that this result also applies to the infinite-horizon model.)

2.3 Output variability

Of course, to move prices back towards their initial level following a shock will likely require larger deviations in real interest rates and, consequently, in output, than would be necessary if the shock to prices were accommodated. For example, in the case illustrated above, where $\lambda = 0$, price-level targeting requires that $BZ_{2|1} = -\frac{\varepsilon_1}{2}$, whereas inflation targeting requires that $BZ_{2|1} = 0$. If one supposes that

$$BZ_t \equiv by_t + CX_t, \quad (14)$$

where X_t is a vector of variables orthogonal to ε_t , then price-level targeting increases the standard deviation of output at time 2 by $\frac{\sigma_\varepsilon}{2b}$ over that under inflation targeting. Thus, we can make the following conclusion:

*Under plausible specifications, price-level targeting is likely to generate higher variability in output than inflation targeting.*¹⁰

When output variability is a concern for welfare (i.e., $\lambda \neq 0$), one must weigh the marginal benefit of moving prices back towards their initial level against the marginal cost of larger output variations. From the perspective of overall social welfare, the marginal benefit of (an anticipated) monetary action at time 2 equals

$$\frac{\partial}{\partial i_{2|1}}(\pi_1^2 + \pi_2^2) = 2(\pi_1 + \pi_2)\frac{\partial \pi_2}{\partial i_{2|1}} = 2p_2\frac{\partial \pi_2}{\partial i_{2|1}}, \quad (15)$$

10. Assuming that an equal weight, λ , is placed on output variability in both regimes.

where we have used the fact that inflation in period 1 depends one-for-one on expected inflation in period 2, and the price-level at time 0 is assumed to equal 0. But this is precisely the marginal benefit, from the standpoint of a price-level targeter, of a monetary action at time 2, given the price level at time 1; e.g.,

$$\frac{\partial}{\partial i_{2|1}}(p_2^2) = 2p_2 \frac{\partial p_2}{\partial i_{2|1}} = 2p_2 \frac{\partial \pi_2}{\partial i_{2|1}}. \quad (16)$$

Thus, whether from the standpoint of overall social welfare or that of a price-level targeter, the marginal benefit of (an anticipated) monetary action at time 2 is the same.

On the other hand, the marginal costs of a monetary action at time 2 from the standpoint of overall social welfare and from that of a price-level targeter (or an inflation targeter) are, respectively,

$$\frac{\partial}{\partial i_{2|1}}(\lambda y_1^2 + \lambda y_2^2) = 2\lambda y_1 \frac{\partial y_1}{\partial i_{2|1}} + 2\lambda y_2 \frac{\partial y_2}{\partial i_{2|1}} \quad (17)$$

and

$$\frac{\partial}{\partial i_{2|1}}(\lambda y_2^2) = 2\lambda y_2 \frac{\partial y_2}{\partial i_{2|1}}. \quad (18)$$

Thus, if output in period 1, y_1 , is not affected by monetary actions in the future (e.g., if output is backward-looking), then $\frac{\partial y_1}{\partial i_{2|1}} = 0$, and the marginal costs are equal under the two regimes.

Since the marginal benefits are also equal, it follows that price-level targeting coincides with the first-best policy under these conditions, and therefore dominates inflation targeting in terms of overall welfare.¹¹ This applies even though output may be more variable under price-level targeting.¹²

If output in period 1 is actually affected by monetary actions in the future (e.g., if output is forward-looking), then monetary action at time 2 would entail an output deviation at time 1 that is ignored under price-level targeting. In this case, the two policies are not likely to coincide. Nonetheless, the first-best policy is still likely to move prices back towards their initial level; e.g.,

11. When the same weight, λ , is placed on output variability in both regimes.

12. In a very simplified model, Vestin (1999) shows that for any weight λ , a weight λ' can be found so that price-level targeting, defined with the weight λ' placed on output variability instead of λ , outperforms inflation targeting in terms of both inflation and output variability.

assuming that ε_1 is non-negatively correlated with output in period 1,¹³ $\frac{\partial y_1}{\partial i_{2|1}} \leq 0$, and that deflation in period 2 can be achieved only through output contraction in period 2, then one can verify that a deflation in period 2, following inflation in period 1, would be welfare-enhancing. That is, the overall marginal welfare effect of a deflation in period 2 is positive; e.g., with $\varepsilon_1 > 0$,

$$\frac{\partial L_1}{\partial i_{2|1}} = 2(\pi_1 + \pi_2) \frac{\partial \pi_2}{\partial i_{2|1}} + 2\lambda y_1 \frac{\partial y_1}{\partial i_{2|1}} + 2\lambda y_2 \frac{\partial y_2}{\partial i_{2|1}} \leq 0 \text{ at } y_2 = 0. \quad (19)$$

In fact, under some circumstances, the first-best policy may bring prices closer to their initial level in period 2 than price-level targeting would. Consider, for instance, a scenario where, as a direct consequence of a shock, output is positive in period 1 but unaffected in period 2. Then, a contraction in period 2 has the benefit of constraining output as well as inflation in period 1; thus, from the standpoint of overall welfare, there is an added incentive to produce a contraction in period 2. It follows that a larger contraction may be called for under the first-best policy than under price-level targeting.

While suggestive, these results are nevertheless particular to the two-period model, even if output is backward-looking. In a three-period horizon model, for example, similar derivations to those employed above show that the marginal benefits of an anticipated monetary action at time 3 continue to be equal whether they are calculated from the standpoint of overall social welfare or from the standpoint of a price-level targeter. But the same may not be true with respect to a monetary action at time 2. If there is no endogenous persistence (e.g., $\frac{dBZ_{3|1}}{di_{2|1}} = 0$), then the marginal benefit of a monetary action at time 2 in terms of overall social welfare and from the standpoint of a price-level targeter is, respectively,¹⁴

13. Which, of course, applies if ε_1 is a demand shock.

14. In general, the marginal benefit of (an anticipated) monetary action at time 2 in terms of overall social welfare equals

$$\frac{\partial}{\partial i_{2|1}}(\pi_1^2 + \pi_2^2 + \pi_3^2) = 2(\pi_1 + \pi_2) \frac{dBZ_{2|1}}{di_{2|1}} + 2(\pi_1 + \pi_2 + \pi_3) \frac{dBZ_{3|1}}{di_{2|1}} = 2p_2 \frac{dBZ_{2|1}}{di_{2|1}} + 2p_3 \frac{dBZ_{3|1}}{di_{2|1}},$$

whereas the marginal benefit of that action from the standpoint of a price-level targeter, taking price p_1 as given, equals

$$\frac{\partial}{\partial i_{2|1}}(p_2^2 + \hat{p}_3^2) = 2(p_2 + \hat{p}_3) \frac{dBZ_{2|1}}{di_{2|1}} + 2(p_2 + 2\hat{p}_3) \frac{dB\hat{Z}_{3|1}}{di_{2|1}},$$

where the hat superscript refers to the optimal value of a variable at time 3.

$$\frac{\partial}{\partial i_{2|1}}(\pi_1^2 + \pi_2^2 + \pi_3^2) = 2(\pi_1 + \pi_2)\frac{dBZ_{2|1}}{di_{2|1}} = 2p_2\frac{dBZ_{2|1}}{di_{2|1}} \quad (20)$$

and

$$\frac{\partial}{\partial i_{2|1}}(p_2^2 + \hat{p}_3^2) = 2(p_2 + \hat{p}_3)\frac{dBZ_{2|1}}{di_{2|1}}, \quad (21)$$

where \hat{p}_3 refers to the price level that applies at time 3 under price-level targeting (once all variables at time 2 have been determined). As equations (20) and (21) show, except in the case where $\hat{p}_3 = 0$, which applies when $\lambda = 0$ (see section 2.5), a price-level targeter has a greater incentive to move prices back to their initial level at time 2 than the first-best policy would call for. This is because, from the standpoint of a price-level targeter, the reduction in the price level at time 2 implies a reduction in the price level in all future times, and is therefore witnessed in the objective function of the price-level targeter in all future periods, whereas the reduction has no implications for social welfare in future periods when only inflation stability is the concern.

Endogenous persistence of the effects of monetary actions on future state variables (e.g., $\frac{dBZ_{3|1}}{di_{2|1}} \neq 0$), adds another complication to the problem (see the derivations of the marginal benefit and marginal cost in footnote 14). Exogenous persistence (i.e., serial correlation of the shocks) does not affect the qualitative arguments above,¹⁵ but it can affect the relative merit of the various policies in quantitative terms. In general, the relative merit of price-level targeting and inflation targeting depends on all of the model's parameters, including the data-generating process governing the variables.

It can be shown, for example, that in the three-period horizon model, inflation targeting is more favourable than price-level targeting if equal weight is placed on output and inflation variability in the welfare function. However, if the shock ε_1 is transitory rather than white noise—specifically, if $\varepsilon_2 = -\frac{1}{2}\varepsilon_1$ —then price-level targeting would clearly be optimal. Indeed, under the latter conditions, price-level targeting would call for no monetary action at time 2, since the price level converges back to the target of its own accord, whereas inflation targeting would perversely offset the shock ε_2 and move the price level back towards the level at the end of period 1.

15. To see this, note that one can rewrite the Phillips curve in the form $\pi_t = \pi_{t+1|t} + \hat{B}\hat{Z}_{t|t-1} + v_t$, where $\hat{B}\hat{Z}_t \equiv BZ_t + \varepsilon_{t|t-1}$, $\varepsilon_t \equiv \varepsilon_{t|t-1} + v_t$ and v_t is white noise, and thus revert to the previous case with white-noise shocks.

2.4 The role of forward-looking behaviour

Present theory supports the type of Phillips curve described in equation (1). However, empirical studies seem to support a hybrid equation of the form

$$\pi_t = a\pi_{t+1|t} + (1-a)\pi_{t-1} + CX_{t|t-1} + \varepsilon_t \quad (a < 1), \quad (22)$$

whereby current inflation is partly influenced by future expectations and partly by past inflation. Since forward-looking expectations by agents are at the heart of the arguments presented above, it is important to examine how alternative specifications such as equation (15) may alter the conclusions of the previous sections. For that purpose, suppose now that the Phillips curve has the form

$$\pi_t = a\pi_{t+1|t} + BZ_{t|t-1} + \varepsilon_t, \quad (23)$$

where $0 \leq a \leq 1$.

Assuming again that $\lambda = 0$, for simplicity, elementary calculus shows that, under the *first-best policy*,

$$BZ_{2|1} = \pi_{2|1} = -a\pi_1 \quad (24)$$

$$\pi_1 = \frac{1}{1+a^2}\varepsilon_1 \quad (25)$$

$$L_1 = \frac{1}{1+a^2}\sigma_\varepsilon^2. \quad (26)$$

Thus, the smaller is a (i.e., the less prices are influenced by future expectations), the smaller is the capacity of expected future monetary actions to restrain current inflation, and the larger is the social loss. In an extreme case, where $a = 1$, one reproduces the result of the earlier section that the first-best policy consists of returning prices to their initial level: $\pi_{2|1} = -\pi_1$. In another extreme case, where $a = 0$, the central bank cannot offset contemporaneous inflationary shocks by means of expectations. Consequently, it accommodates shifts in the price level and targets 0 inflation at time 2, as would apply under inflation targeting. In general, with $0 < a < 1$, the first-best policy involves a partial return of prices to their initial level, as witnessed by equation (24).

Inflation targeting calls for 0 inflation at time 2. The outcome under this policy is therefore independent of a , and is identical to the outcome that applies under the first-best policy when $a = 0$; e.g.,

$$\pi_1 = \varepsilon_1 \quad (27)$$

$$BZ_{2|1} = \pi_{2|1} = 0 \quad (28)$$

$$L_1 = \sigma_\varepsilon^2. \quad (29)$$

Price-level targeting calls for returning prices at time 2 to their initial level. Therefore, this policy is also independent of a , and is identical to the first-best policy that applies when $a = 1$:

$$BZ_{2|1} = \pi_{2|1} = -\pi_1. \quad (30)$$

$$\pi_1 = \frac{1}{1+a} \varepsilon_1 \quad (31)$$

$$L_1 = \frac{2}{(1+a)^2} \sigma_\varepsilon^2. \quad (32)$$

Thus, as witnessed by the magnitude of the loss L_1 in each case, we can make the following conclusion:

*The more forward-looking the prices (i.e., the larger is the coefficient a), the greater is the benefit of price-level targeting over inflation targeting, and the closer it is to the first-best policy.*¹⁶

2.5 The infinite-horizon model

Besides abstracting from certain complications that arise in a longer horizon model, as discussed in section 2.3, the two-period model specification does not reveal unambiguously the trend of prices in the long run under the alternative policies. Thus, while the previous results may show that, under the first-best policy, prices partially return to their initial level in the second period, it is not clear whether prices would eventually fully return to their initial level if the horizon were infinite.¹⁷

16. Price-level targeting outperforms inflation targeting if $a > \sqrt{2} - 1$, in the case $\lambda = 0$.

17. Some authors consider the convergence of prices to a target in the long run, rather than optimality with respect to the objective function in (3), as the defining condition of price-level targeting (Barnett and Engineer 2000).

Suppose, therefore, that there are infinitely many periods, and that the social welfare function has the form:

$$-L_t = -E_t \left[\sum_{i=0}^{\infty} \beta^i (\pi_{t+i}^2 + \lambda y_{t+i}^2) \right]. \quad (33)$$

Define the various types of policies as in the two-period case, and consider a single shock, ε_1 , at time 1.

Then, except in odd cases, inflation targeting would not return prices to the target, since, in every period, once by-gones are by-gones, the incentive is to cause no further changes in prices. Price-level targeting, on the other hand, would always return prices to the target, since, in each period, there is an incentive to bring prices closer to the target.

As to the first-best policy, as of time 1, the overall cost of moving prices eventually back towards their initial level consists of the (squared) deviations in inflation induced in the subsequent periods plus the (squared) deviations in output necessary for that purpose. It follows that, if output variability is not a concern to welfare (i.e., $\lambda = 0$), then inducing movements in prices beyond period 2 cannot be efficient, for inflation in period 1 is affected only through expected inflation in period 2. Thus, we can make the following conclusion:

If output variability does not affect welfare (i.e., $\lambda = 0$), then the first-best policy in the infinite-horizon model is identical to that in the two-period model. In particular, in this case, prices converge back to their initial level under the first-best policy if, and only if, the coefficient, a , on inflation expectations (in equation (18)) equals 1.

In general, as noted in section 2.3, when output variability does affect welfare (i.e., $\lambda \neq 0$), then the first-best policy in the infinite-horizon case depends on the full data-generating process of all the variables in the model, and does not coincide with price-level targeting, even when the Phillips curve is fully forward-looking.¹⁸

Nonetheless, it can be shown that, typically, when the Phillips curve is fully forward-looking, the first-best policy does eventually return prices to their initial level. To illustrate a very simple case, suppose that the vector of explanatory variables, Z_t , in the Phillips curve consists solely of the

18. Using the simple model described below, except that the interest rate can affect inflation contemporaneously, Vestin (1999) shows that, when shocks are white noise, for any given weight λ , one can find a different weight λ' , so that the first-best policy based on λ is identical to price-level targeting based on λ' ; when the shocks are persistent, λ' does not exist.

output gap, and that the latter is negatively proportional to the real interest rate, so that we can think of the output gap as the instrument of policy.

Define the Lagrangian,

$$\min E_1 \left\{ \sum_{t=1}^{\infty} \beta^t [(\pi_t^2 + \lambda y_t^2) + 2\phi_t(\beta\pi_{t+1} + By_t + \varepsilon_t - \pi_t)] \right\}. \quad (34)$$

The first-order conditions then give

$$\pi_t = \phi_t - \phi_{t-1}, \text{ for } t > 1 \quad (35)$$

$$\pi_1 = \phi_1, \quad (36)$$

$$\lambda y_t = -\phi_t B, \text{ for } t > 1. \quad (37)$$

Combining the three equations, it follows that

$$p_T - p_0 = \sum_{t=1}^T \pi_t = \phi_T = -\frac{\lambda}{B} y_T. \quad (38)$$

Hence, since the output gap must converge to 0 in the long run, the price level returns to its level at time 0.

3. Small Open Economy

This section extends the analysis to the case of a small open economy with a flexible exchange rate, such as Canada's economy. We also briefly discuss whether it is practical for Canada to pursue a policy that is very different from the policy pursued by its major trading partner, the United States.

Suppose that the transmission mechanism in the open economy is represented by the following three equations:

$$\begin{aligned} \pi_t &= a\pi_{t+1|t} + (1-a)\pi_{t-1} + by_{t|t-1} - f(e_{t|t-1} - e_{t-1}) + \Phi X_{t|t-1} + \varepsilon_t \\ y_t &= -cr_{t|t-1} - ge_{t|t-1} + \Psi X_{t|t-1} + \eta_t \\ e_t &= hr_t + \Omega X_t + v_t, \end{aligned} \quad (39)$$

where r_t is the real interest rate, defined as the nominal interest rate minus one-period-ahead expected inflation, $\pi_{t+1|t}$; e_t is the real exchange rate, defined as $e_t \equiv p_t + s_t - p_t^f$, where s_t is

the nominal exchange rate (e.g., the price of a unit of domestic currency in terms of foreign currency), p_t is the domestic price level, and p_t^f is the foreign counterpart; X_t is a vector of exogenous variables, such as U.S. variables and real commodity prices; ε_t , η_t , and v_t are white-noise shocks; and, as before, π_t is the inflation rate and y_t is the output gap. Again, all variables equal 0 at steady state, and the economy returns automatically to the steady state in period 3. The central bank sets the interest rate at time t before observing the shock ε_t .

The first two equations stand for a Phillips curve and an IS curve, respectively, and the third equation describes the link between the interest rate and the exchange rate. The exchange rate term in the Phillips curve witnesses the direct effect of the exchange rate on inflation—for instance, through its effect on import prices—while the exchange rate term in the IS curve witnesses the effect of that variable on demand through its effect on net exports. The lag structure assumed in the model is convenient, but not essential for our purposes. The crucial elements are the dependence of current inflation, π_t , on expected future inflation, $\pi_{t+1|t}$, and the inability of monetary actions to immediately offset unanticipated shocks to inflation.¹⁹

Although the exchange rate introduces a second channel through which monetary policy can affect the economy in this model, that fact does not alter the conclusions obtained earlier in the closed-economy case. To see this formally, set

$$BZ_t \equiv (1 - a)\pi_{t-1} + by_t - f(e_t - e_{t-1}) + \Phi X_t \quad (40)$$

to transform the Phillips curve above into the form (1) considered in the previous section, and notice that, since

$$\begin{aligned} BZ_{t|t-1} &= (1 - a)\pi_{t-1} + by_{t|t-1} - f(e_{t|t-1} - e_{t-1}) + \Phi X_{t|t-1} \\ &= (1 - a)\pi_{t-1} - (bc + fh + bgh)r_{t|t-1} + (b\Psi - f\Omega - bg\Omega + \Phi)X_{t|t-1} + fe_{t-1}, \end{aligned} \quad (41)$$

$r_{t|t-1}$ determines the value of $BZ_{t|t-1}$.

Thus, in principle, as long as the welfare function has the same form as in the previous section, the conclusions obtained earlier on the relative merits of price-level targeting ought to apply equally to the small open economy represented above. Moreover, one may argue that this benefit is independent of the monetary regime adopted by a trading partner, whether, for instance, the trading partner is itself targeting the price level or the inflation rate.

19. Alternatively, one obtains similar results if the cost in terms of output variability of offsetting shocks contemporaneously is very high.

There are, however, a number of important qualifications to this conclusion. It is useful to examine first one argument in support of the claim that the monetary regime adopted by the trading partner is irrelevant. Namely, it is argued that a flexible exchange rate would automatically adjust to alternative monetary regimes in the foreign country to keep the real exchange rate, and therefore the choice of domestic policy, unaffected. This is true if the alternative regimes under consideration in the foreign country differ only in nominal terms, and the difference is reflected solely in the behaviour of the nominal exchange rate. In other words, the alternative regimes are supposed to induce identical behaviour of the exogenous variables, X_t . Under these conditions, indeed, the regime that is adopted in the foreign country should not have any bearing on the choice of domestic policy. For example, everything else being equal, the trend inflation rate targeted by the foreign country, whether it is 2 per cent or 3 per cent, should have no bearing on the domestic economy and, a fortiori, its monetary policy.

If, however, the alternative regimes under consideration in the foreign country differ in other than just nominal terms, as would be the case, for instance, if the alternative regimes are inflation and price-level targeting, then it is most likely that exchange rate adjustments would not completely insulate the domestic economy from the changes in behaviour in the foreign country. In other words, the differences in behaviour of the alternative regimes will be reflected in differences in the data-generating process governing the exogenous variables, X_t , that enter the model. Under these conditions, the regime adopted in the foreign country would certainly affect the choice of domestic policy insofar as that choice depends on the data-generating process governing all the variables in the model.²⁰

In fact, it is not at all clear how the exchange rate would adjust to different behaviours in prices or output in the foreign country. Movements in the exchange rate are quite uncertain, in general, and this uncertainty is likely to be larger, the larger the spread between domestic and foreign interest rates—i.e., the more divergent the monetary policies are in the two countries. One possible reason for this fact is that the foreign country has private information about the state of its economy and/or the path that its policy will take in the future in response to current events. By imitating its policy, the domestic country can reduce the risk of error regarding conditions in the foreign country. Another possible reason is that the exchange rate will be better anchored if the spread in interest rates between the two countries remains close to the equilibrium level. A third possible reason is

20. There is some indication that the above criticism may not be too serious. Namely, section 2 showed that in the two-period-horizon model, when the Phillips curve is fully forward-looking, price-level targeting coincides with the optimal policy, irrespective of the behaviour of the exogenous shocks. Section 2 further showed that the solutions to the infinite-horizon model are identical to those in the two-period model when output variability is not a concern for welfare. These results suggest that the relative merits of price-level targeting may be robust to the behaviour of exogenous shocks.

that, while one might reasonably expect the exchange rate elasticity to various shocks to remain stable if the state conditions are the same in both countries and the two economies pursue identical policies, this is not necessarily the case if the two countries pursue very different monetary policies.

To the extent that movements in the exchange rate affect both output and inflation, there is an incentive to limit the uncertainty about such movements, and therefore to pursue a domestic policy that is somewhat similar to the policy pursued by the trading partner.

Another qualification is that the welfare function is presumed to depend solely on output and inflation variability, and ignores the costs of exchange rate variability per se. To the extent that exchange rate variability (in its own right) is detrimental to the economy, and price-level targeting induces a more variable nominal exchange rate, the latter policy will be less advantageous than inflation targeting. Admittedly, aside from direct costs of exchange rate transactions, empirical studies so far have not been able to find conclusive evidence regarding the welfare costs of exchange rate variability,²¹ but the same can be said regarding the welfare implications of price variability. At an intuitive level, one would expect nominal exchange rate variability to be detrimental to foreign investors and importers/exporters for the same reasons that price variability is detrimental to domestic consumers, although the former are likely to be capable of hedging against risks at a lower cost.

Finally, at a more fundamental level, the model does not distinguish between different sectors of the economy—specifically, the traded and non-traded sectors—and the effect of policy on the terms of trade and the current account balance.

4. Conclusion

This paper has reviewed the argument in favour of price-level targeting in the context of a closed economy. We have shown that price-level targeting is more advantageous the more forward-looking the price-setting mechanism in the economy, and the less weight placed on output variability in the welfare function. In general, the first-best policy involves only a partial return of prices to their initial level following shocks. The extent and the rate at which the first-best policy moves prices back towards their initial level following shocks—hence the relative merits of price-level targeting—depend on all of the model's specifications, and are empirical questions.

21. See, for instance, Coté (1994) for a survey of the literature on exchange rate volatility and trade, and Lafrance and Tessier (2000) for a study of the relationship between exchange rate variability and investment in Canada.

The paper has also examined the case of a small open economy with a flexible exchange rate, such as Canada's economy. We have shown that, in the context of a simple model, the results extend to the open economy case. However, we have argued that this conclusion may be too simplistic for a country such as Canada, and that it ignores important factors, such as the cost of the exchange rate volatility that would ensue if Canada and the United States pursued very dissimilar policies, and the effects on the terms of trade. The latter discussion was mostly descriptive. Further research is needed, especially within a general-equilibrium framework, to provide more satisfactory answers.

References

- Barnett, R. and M. Engineer. 2000. "When Is Price-Level Targeting a Good Idea?" In *Price Stability and the Long-Run Target for Monetary Policy*, 101–43. Proceedings of a seminar held by the Bank of Canada, June. Ottawa: Bank of Canada.
- Clarida, R., L. Gali, and M. Gertler. 1999. "The Science of Monetary Policy: A New Keynesian Perspective." *Journal of Economic Literature* 37 (4): 1661-1707.
- Coté, A. 1994. "Exchange Rate Volatility and Trade: A Survey." Bank of Canada Working Paper No. 94-5.
- Dittmar, R. and W.T. Gavin. 2000. "What Do New-Keynesian Phillips Curves Imply for Price-Level Targeting?" *Federal Reserve Bank of Saint Louis Review* 82 (2): 21-30.
- Duguay, P. 1994. "Some Thoughts on Price Stability versus Zero Inflation." Paper presented at a conference on Central Bank Independence and Accountability, Università Bocconi, Milan.
- Lafrance, R. and D. Tessier. 2000. "Exchange Rate Variability and Investment in Canada." Manuscript.
- McCallum, B. 1997. "Issues in the Design of Monetary Policy Rules." NBER Working Paper No. 6016.
- Rogoff, K. 1985. "The Optimal Degree of Commitment to an Intermediate Monetary Target." *Quarterly Journal of Economics* 100 (November): 1169-89.
- Svensson, L. 1999. "Price Level Targeting versus Inflation Targeting: A Free Lunch?" *Journal of Money, Credit and Banking* 31: 277-95.
- Svensson, L. and M. Woodford. 1999. "Implementing Optimal Policy through Inflation-Forecast Targeting." Manuscript.
- Vestin, D. 1999. "Price Level Targeting versus Inflation Targeting in a Forward Looking Model." Working Paper Series ISSN 1402-1903, No. 106, Sveriges Riksbank, Stockholm.
- Woodford, M. 1999. "Optimal Monetary Policy Inertia." NBER Working Paper No. 7261.

Bank of Canada Working Papers

Documents de travail de la Banque du Canada

Working papers are generally published in the language of the author, with an abstract in both official languages. *Les documents de travail sont publiés généralement dans la langue utilisée par les auteurs; ils sont cependant précédés d'un résumé bilingue.*

2001

2001-23	Modelling Mortgage Rate Changes with a Smooth Transition Error-Correction Model	Y. Liu
2001-22	On Inflation and the Persistence of Shocks to Output	M. Kichian and R. Luger
2001-21	A Consistent Bootstrap Test for Conditional Density Functions with Time-Dependent Data	F. Li and G. Tkacz
2001-20	The Resolution of International Financial Crises: Private Finance and Public Funds	A. Haldane and M. Kruger
2001-19	Employment Effects of Restructuring in the Public Sector in North America	P. Fenton, I. Ip, and G. Wright
2001-18	Evaluating Factor Models: An Application to Forecasting Inflation in Canada	M.-A. Gosselin and G. Tkacz
2001-17	Why Do Central Banks Smooth Interest Rates?	G. Srouf
2001-16	Implications of Uncertainty about Long-Run Inflation and the Price Level	G. Stuber
2001-15	Affine Term-Structure Models: Theory and Implementation	D.J. Bolder
2001-14	L'effet de la richesse sur la consommation aux États-Unis	Y. Desnoyers
2001-13	Predetermined Prices and the Persistent Effects of Money on Output	M.B. Devereux and J. Yetman
2001-12	Evaluating Linear and Non-Linear Time-Varying Forecast-Combination Methods	F. Li and G. Tkacz
2001-11	Gaining Credibility for Inflation Targets	J. Yetman
2001-10	The Future Prospects for National Financial Markets and Trading Centres	C. Gaa, S. Lumpkin, R. Ogrodnik, and P. Thurlow
2001-9	Testing for a Structural Break in the Volatility of Real GDP Growth in Canada	A. Debs
2001-8	How Rigid Are Nominal-Wage Rates?	A. Crawford

Copies and a complete list of working papers are available from:

Pour obtenir des exemplaires et une liste complète des documents de travail, prière de s'adresser à :

Publications Distribution, Bank of Canada
234 Wellington Street, Ottawa, Ontario K1A 0G9
E-mail: publications@bankofcanada.ca
Web site: <http://www.bankofcanada.ca>

Diffusion des publications, Banque du Canada
234, rue Wellington, Ottawa (Ontario) K1A 0G9
Adresse électronique : publications@banqueducanada.ca
Site Web : <http://www.banqueducanada.ca>