Do Inflation-Targeting Central Banks Implicitly Target the Price Level?*

Francisco Ruge-Murcia
University of Montréal

This paper examines the time-series properties of the price level in five inflation-targeting countries. For the regimes in Australia, New Zealand, Sweden, and the United Kingdom, the price level wanders away from the path implied by the inflation target, and test results suggest that it has a unit root. For the regime in Canada, the price level tracks the path implied by the target and test results partly support the view that it is covariance stationary. These results do not mean that Canada covertly follows a price-level targeting regime but suggest, instead, heterogeneity in the actual application of inflation targeting across countries. Survey data on inflation expectations are inconclusive as to whether agents treat the Canadian price level as a trend-reverting process.

JEL Codes: E3, E5.

1. The Debate

A current and important debate in monetary policy concerns the relative advantages of price-level targeting compared with inflation

*This paper was inspired by a question by Angelo Melino during the Bank of Canada Economic Conference held in Ottawa in November 2009, and its first draft was written while I was visiting the International Department of the Bank of Canada. I thank the International Department for its hospitality; Guido Ascari, Mick Devereux, and Andrew Levin for comments and suggestions; Julie Arbour and Andrew Levin (again) for providing me with survey data on inflation expectations; and the Social Sciences and Humanities Research Council of Canada for financial support. The views expressed here are those of the author and should not be attributed to either the Bank of Canada or its staff. The data used in this analysis are available at my web page: https://sites.google.com/site/frugemurcia. Author contact: Department of Economics, University of Montréal, C.P. 6128, succursale Centre-ville, Montréal (Québec) H3C 3J7, Canada. E-mail: francisco.ruge-murcia@umontreal.ca.

targeting. Under inflation targeting, the central bank announces a specific quantitative target for the inflation rate and directs monetary policy towards that goal. The policy specifies the measure of inflation to be targeted, whether the target is a single value or a range, and the horizon at which the target applies. Among the countries that currently target inflation are Canada, Sweden, and the United Kingdom. Under price-level targeting, the central bank targets the path of a price index. This path may involve an upward trend implying a positive inflation rate. The debate is of practical importance because some successful inflation-targeting central banks have recently considered the possibility of adopting a price-level target instead.[1]

An obvious difference between inflation targeting and price-level targeting is that the former targets the rate of growth of the price level while the latter targets the price level itself. This means that under price-level targeting, an increase in the price level that is larger than allowed by the policy has to be offset in subsequent periods. In theory, inflation targeting imposes no such obligation on the central bank, and, hence, temporary shocks that impinge on the price level may be accommodated. Put differently, under inflation targeting "bygones are bygones." A simple characterization of both policies is plotted in columns 1 and 2 of figure 1 (column 3 will be discussed later). The top panels plot the path of the targeted variable, whether inflation or the price level. The lower panels plot the implied paths for the price level (in the case of inflation targeting) or the inflation rate (in the case of price-level targeting). In this illustration, the inflation target is 2 percent and the price-level target involves an upward trend of 2 percent.

Consider first inflation targeting in column 1. Imagine that a temporary shock takes inflation from 2 percent in period 1 to, say, 3 percent in period 2 (panel A). Then, in period 3, the central bank steers inflation back to the targeted value of 2 percent and manages to keep it at this value thereafter. This inflation-targeting policy would be considered successful in that inflation was returned to the target and the (asymptotic) mean of inflation is the target value. The path for the price level implied by this policy (panel B) shows that

[1]See "Renewal of the Inflation-Control Target: Background Information," issued by the Bank of Canada in November 2006.

Figure 1. Inflation versus Price-Level Targeting

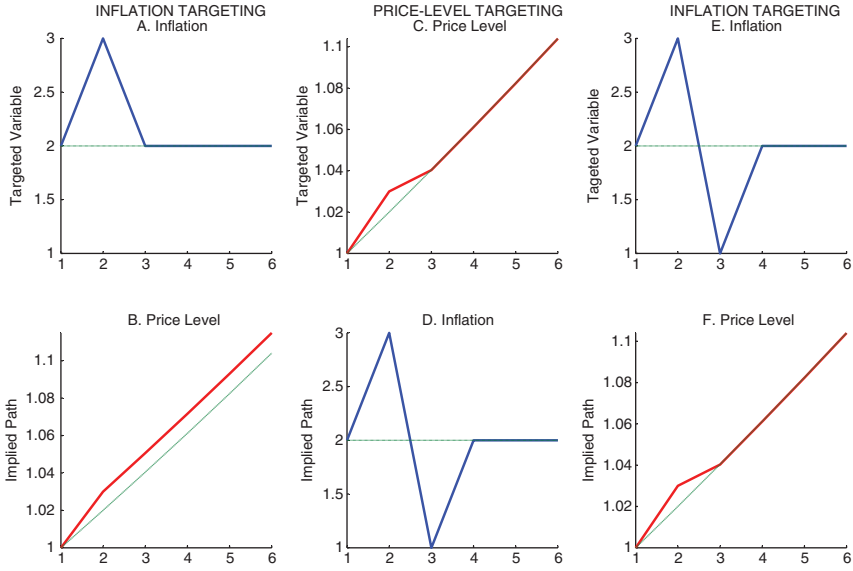

the price index grows by 3 percent between periods 1 and 2, and by 2 percent thereafter. Notice that the temporary shock has a permanent effect on the price level: since the central bank accommodates the shock, bygones are bygones.

Consider now the same shock but under a price-level targeting policy. In column 2 (panel C), the price level grows by 3 percent between periods 1 and 2, and the price index is, therefore, higher than the value targeted for that period. In order to return to the desired path, the central bank must set monetary policy so that the price level grows by only 1 percent from period 2 to 3, and by 2 percent thereafter. This price-level targeting policy would be considered successful in that the price index was returned to the targeted path. In this case, the temporary shock does not have a permanent effect on the price level. In panel D, note that after an inflation rate of 3 percent in period 1, the policy induces an inflation rate of 1 percent in period 3 and of 2 percent thereafter.

Finally, let us consider a slightly different characterization of inflation targeting. In column 3 (panel E), following the temporary shock that takes inflation to 3 percent in period 2, the central bank

sets monetary policy so that inflation is 1 percent in period 3 and 2 percent thereafter. This inflation-targeting policy would be considered successful in that inflation was returned to the target and the mean of inflation is the target value. In some sense this version of inflation targeting may be regarded as more successful than the one in column 1 because average inflation is the target value in the short run, rather than only asymptotically. Comparing columns 2 and 3, it is clear that this policy delivers a price-level path that is identical to the one obtained under price-level targeting. In particular, temporary shocks have no permanent effect on the price level. Hence, there is an observational equivalence between inflation and price-level targeting. In this figure, the observational equivalence arises from the purposeful policy action of the inflation-targeting central bank, which seeks to deliver average inflation rates close to the target rate in the short run. In principle, this equivalence may also arise as a result of symmetric shocks that take inflation sometimes above, sometimes below, its target.

Earlier literature finds that optimal monetary policy under commitment involves a stationary price level. This is so in New Keynesian models (see Clarida, Gali, and Gertler 1999 and Woodford 1999) and in models with rational inattention (Ball, Mankiw, and Reis 2005). Since price-level targeting induces stationarity in the price level while inflation targeting does not, this is a powerful argument in favor of the former policy. In models where commitment is not possible, price-level targeting has a number of desirable implications which are welfare improving compared with inflation targeting. Most of these implications stem from the fact that forward-looking agents incorporate the future price-level path into their current actions. Svensson (1999) shows that price-level targeting reduces the variance of inflation and eliminates any average inflation bias. Vestin (2006) finds that even if social preferences concern inflation, welfare is increased by delegating monetary policy to a central banker with an explicit price-targeting objective. Eggertsson and Woodford (2003) and Wolman (2005) find that price-level targeting is helpful in overcoming distortions in the neighborhood of the zero lower bound on nominal interest rates, and Amano, Ambler, and Ireland (2007) show that the optimal level of wage indexation is lower under price-level targeting and the resulting real wage flexibility is welfare improving. On the other hand, in a setup where some of the agents

are backward looking, price-level targeting is not optimal and may be dominated by a policy that allows some price-level drift (Steinsson 2003).

This paper contributes to the debate by examining the time-series properties of the price level in five inflation-targeting countries. Data plots show that in Australia, New Zealand, Sweden, and the United Kingdom the price level wanders away from the path implied by the inflation target, as in a bygones-are-bygones inflation-targeting regime, while in Canada the price level closely tracks this path, as in a price-level targeting regime. The predictions of bygones-are-bygones inflation targeting for the price level are tested against those of price-level targeting using unit-root tests. Subject to the caveat that it is difficult to differentiate a stationary process from a unit-root process in a finite sample, results suggest that for Canada (but not for the other countries) the predictions of inflation targeting are rejected in favor of those of price-level targeting. These results do not mean that Canada covertly pursues a price-level targeting policy. Instead, they suggest heterogeneity in the application of inflation targeting across countries, with the Canadian regime seemingly able to deliver a predictable price level.

On the question of whether Canadian market participants treat the price level as trend reverting, I find that a forecasting model that assumes trend reversion has lower root mean squared error than a model that assumes that price-level shocks are permanent. However, survey data on expectations provide only suggestive evidence that agents forecast the price level as a trend-reverting process. Thus, an explicit price-level targeting policy may be welfare improving if it helps stabilize price-level expectations around a pre-announced path.

2. The Data

2.1 Five Inflation-Targeting Countries

The empirical analysis is based on price-level indices from Australia, Canada, New Zealand, Sweden, and the United Kingdom. In order to motivate the choice of indices and sample periods used in the analysis, I briefly review the specific institutional arrangements in each country.

In Australia, the Reserve Bank of Australia (RBA) has had an inflation target since early 1993.[2] Since 1996 the policy has been formalized in a Statement on the Conduct of Monetary Policy jointly issued by the government and the RBA. The target is a range of 2 to 3 percent on average "over the cycle," and initially applied to the growth rate of an "underlying" Consumer Price Index (CPI) that excluded volatile items (for example, gasoline, vegetables, and fruits), public-sector goods, and mortgage interests (see Bernanke et al. 1999, p. 225). Starting in September 1998, the Australian Bureau of Statistics removed interest charges from the CPI, and soon after, the Federal Treasurer and the RBA agreed that applying the target to the CPI was consistent with the intent of the original Statement (see *Reserve Bank Bulletin,* October 1998, pp. 1–5). Thus, for Australia, I use the All Groups CPI taken from the website of the Australian Bureau of Statistics (www.abs.gov.au) for the period 1993:Q1 to 2013:Q1.

In Canada, the Bank of Canada announced "inflation-reduction" targets in February 1991. These targets initially envisaged a reduction in the annual inflation rate to 3 percent by the end of 1992, 2.5 percent by mid-1994, and 2 percent by the end of 1995. As in Australia, the inflation-targeting policy is formalized in a Joint Statement issued by the government and the central bank. Since January 1996, the inflation target is 2 percent, with a tolerance range of plus or minus 1 percentage point. The target is defined in terms of the Consumer Price Index, but a "core" CPI, which excludes food and energy prices and is adjusted for the effects of indirect tax changes, is used as the operational guide for the policy (see the Joint Statement released in February 1998). More recent Statements issued in May 2001 and November 2006 do not make this distinction between indices, reaffirm that the target applies to CPI inflation, and, in fact, do not mention the Core CPI at all.[3] Thus, for

[2]Bernanke et al. (1999, p. 220) date the beginning of inflation targeting in Australia as September 1994, when Governor Fraser stated that "[an] underlying inflation of 2 to 3 per cent is a reasonable goal for monetary policy." The conclusions reported here are robust to using either date.

[3]On the other hand, the background information on the renewal of the inflation targets (see the press releases on May 18, 2001 and November 27, 2006 available at www.bankofcanada.ca/en/press/index.html) still refers to the Core CPI as the operational guide to policy.

Canada, I use the CPI taken from the website of Statistics Canada (www.statcan.gc.ca) for the period 1996:M1 to 2013:M3. I also examine the robustness of results to using the Core CPI (taken from the same source). The sample excludes the transition period to the new inflation-targeting policy, that is, the period of inflation-reduction targets between February 1991 and December 1995.

In New Zealand, the first Policy Target Agreement (PTA) between the Ministry of Finance and the Reserve Bank of New Zealand (RBNZ) was issued in March 1990 and specified a target range of 0 to 2 percent for a (new) Consumer Price Index where housing costs were to be measured as imputed rents rather than as actual expenditures. The PTA of December 1990 respecified the target in terms of the All Groups CPI, and successive Target Agreements have modified the target range and inflation measure. In particular, the PTA of December 1996 increased the width of the range to between 0 and 3 percent, the PTA of December 1997 specified the target in terms of the All Groups CPI excluding credit services, and that of December 1999 specified the target in terms of the All Groups CPI, again. Finally, the PTA of September 2002 adjusted the target to between 1 and 3 percent. Thus, for New Zealand, I use the All Groups CPI taken from the website of the Reserve Bank of New Zealand (www.rbnz.govt.nz) for the period 1991:Q1 to 2013:Q1.[4]

In Sweden, the inflation target was announced in January 1993, but the policy was to take effect only from January 1995 onwards. The target is 2 percent and applies to the annual change in the Consumer Price Index. Until June 2010, the target specified a tolerance range of 1 percentage point above and below. I use the CPI taken from the website of Statistics Sweden (www.scb.se) for the period 1995:M1 to 2013:M3. The sample excludes the transition period to the new inflation-targeting policy, that is, the period from January

[4]The sample excludes the period between 1990:Q2 and 1990:Q4 because I was unable to find the price index used to compute the inflation measure to be targeted under the PTA of March 1990. Part of the problem is that this index was not among those traditionally produced by Statistics New Zealand and, in fact, the PTA explicitly instructed the RBNZ to produce it. However, since the PTA of December of the same year redefined the target in terms of the CPI, which is readily available, and the sample is reduced by only three observations, it seems unlikely that results are affected by the way I have treated the missing observations.

1993 to December 1994 during which the future targeting policy was known but not yet in effect.

Finally, in the United Kingdom, an inflation target range of 1 to 4 percent was announced by the Chancellor in October 1992 and applied (until 2003) to the Retail Price Index excluding mortgage interest payments (RPIX). Starting in June 1995 the target was interpreted as the midpoint of the above range, with thresholds of plus or minus 1.5 percentage points after May 1997 (see Bernanke et al. 1999, p. 154 and 146, respectively). In December 2003, the target was adjusted to 2 percent and applied thereafter to the Consumer Price Index. Thus, for the United Kingdom, I use the RPIX for the period 1992:M10 to 2003:M12 spliced with the appropriately rescaled CPI for the period 2004:M1 to 2013:M3. These indices were taken from the website of the UK National Statistics (www.statistics.gov.uk). I examine the robustness of the results by using the Retail Price Index (RPI) (taken from the same source), rather than the RPIX, in the first part of the sample.

Finally, note that since targets apply to seasonally unadjusted inflation in all countries, the price indices used here are seasonally unadjusted as well.

2.2 Plots

For the five countries in the sample, the upper panels of figure 2 plot the logs of the price index and of the price-level path implied by the inflation target, and the lower panels plot the percentage deviation of the former with respect to the latter. The implied price-level path is computed as

$$\log(P^*_{s+t}) = (t/12)\log(1 + \pi^*/100)\log(P_s), \tag{1}$$

where π^* is the inflation target and s denotes the period when the inflation-targeting policy began. In the cases of Australia and New Zealand, where the data are available at the quarterly frequency and the inflation target is a range, I substitute $t/12$ with $t/4$ in definition (1) and use the midpoint of the range as the empirical counterpart of π^*.

Let us focus first on column 2, which shows that the price level in Canada closely tracks the path implied by the inflation target. The

Figure 2. Price Indices and Deviations from Implied Paths

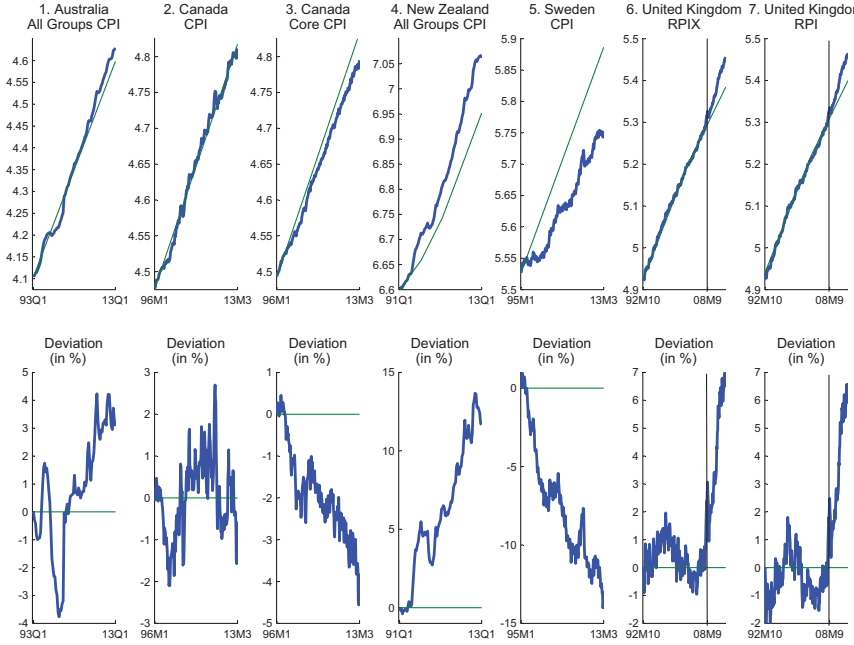

price-level deviations from the implied path are quantitatively small (between −2 and 3 percent) and their sample mean is close to zero (0.006 percent). These observations are at odds with the standard view of inflation targeting, whereby the central bank accommodates shocks to the price level and, as a result, the latter drifts away from the path implied by the inflation target. Instead, the Canadian CPI remains close to the path, just as one would expect to see in a price-level targeting regime. In some sense, this is a real-life analog of the observational equivalence between inflation and price-level targeting illustrated in columns 2 and 3 of figure 1. Columns 1 and 6 show that to some extent this is also true for Australia and the United Kingdom prior to the crisis.[5]

[5]The start of the financial crisis is dated in September 2008, when Lehman Brothers filed for bankruptcy, and it is denoted by a continuous vertical line in the plots for the United Kingdom.

Plots like the upper panel of column 2 (give or take some observations) are also reported by Kamenik et al. (2013) and discussed by Dodge (2005), Parkin (2009), and Melino (2012).[6] Dodge suggests that this observation is due to symmetric shocks and is unlikely to persist in the future. Parkin carefully examines language used in successive Statements and detects a (perhaps) "unconscious" move to price-level targeting (p. 9). Kamenik et al. estimate a reduced-form model of the Canadian economy by Bayesian methods and conclude that there is a low probability that symmetric shocks account for the plot in column 2. Instead, they suggest that an element of price-level targeting in interest rate setting by the Bank of Canada is a more plausible explanation. (These issues are discussed further in the next section.)

Consider now columns 4 and 5 for New Zealand and Sweden, respectively. Note that their plots agree with the usual interpretation of inflation targeting in that the price level drifts away from the path implied by the inflation target. These plots are the real-life analog of the inflation-targeting policy in column 1 of figure 1.

The case of the United Kingdom in column 6 is especially interesting because it illustrates an instance where a large shock (in this case, the financial crisis) may have been accommodated by an inflation-targeting regime. The price level tracks the path implied by the inflation target prior to September 2008 but sharply departs from it thereafter. This observation underscores the inflation "lift-off" in the United Kingdom whereby inflation was well above its target range in the years 2010 and 2011. At the time of writing, it is unclear whether future Bank of England policy will return the price level towards the implied path or whether the observed departure will turn out to be permanent.

Finally, columns 3 and 7 report plots constructed using the Core CPI for Canada and the RPI for the United Kingdom. Column 3 shows that results for Canada are not robust to using the Core CPI rather than the CPI. The difference between columns 2 and 3 is primarily due to the fact that the Core CPI excludes items that account for 17 percent of the CPI and, in the presence of trends in

[6]I am indebted to Bob Amano for bringing these references to my attention.

relative prices, both indexes would have different trends.[7] In contrast, column 7 shows that conclusions for the United Kingdom are robust to using the RPI, instead of the RPIX, in the first part of the sample. The mortgage interest payments that are excluded from the RPI have a weight of only 5 percent and so, compared with the Canadian case, it is more likely that both indices share a common trend.

2.3 Unit-Root Tests

The interpretation of inflation targeting as a policy where bygones are bygones generates two specific and testable predictions about the price level. First, since shocks have a permanent effect, the price level should have a unit root. Second, since subtracting a deterministic trend from a unit-root process does not render the resulting series stationary, then it must be the case that the deviation of the price index from the trend implied by the inflation target should have a unit root as well. In contrast, price-level targeting generates two alternative predictions for the price level. First, since price-level deviations from the targeted path must be offset in future periods, the price level should follow a stationary process around a deterministic trend and, second, the deviations should follow a stationary process around zero.

The natural way to empirically evaluate these implications is by means of unit-root tests. Table 1 reports results of augmented Dickey-Fuller (ADF) unit-root tests for the log of the price level and for the log-deviation of the price level from the path implied by the inflation target. For the former, the hypothesis is that the price level follows a unit-root process with drift and the alternative is a covariance-stationary autoregression with a deterministic trend and a constant term. For the latter, the hypothesis is that the price-level deviation from the implied path follows a unit-root

[7]The Core CPI excludes eight components from the CPI. These components are fruits, vegetables, gasoline, fuel oil, natural gas, mortgage interest, intercity transportation, and tobacco products. The weights of these components in the 2005 basket are 1.25 (fruits, fruit preparations, and nuts), 1.11 (vegetables and vegetable preparations), 4.49 (gasoline), 0.41 (fuel oil and other fuels), 1.35 (natural gas), 5.16 (mortgage interest costs), 1.09 (intercity transportation), and 1.33 (tobacco products and smokers' supplies).

Table 1. ADF Unit-Root Tests

Country and Price Index	Variable					
	Price Level			Deviation from Implied Path		
		Robustness to Augmentation			Robustness to Augmentation	
	Benchmark	6	12	Benchmark	6	12
Australia						
All Groups	−2.244	−2.451	−2.463	−1.013	−1.204	−1.254
Canada						
CPI	−3.514‡	−2.657	−2.101	−3.425‡	−2.555‡	−2.033‡
Core CPI	−2.317	−2.761	−2.544	1.497	1.076	1.289
New Zealand						
All Groups	−2.082	−2.038	−1.883	1.013	1.084	0.974
Sweden						
CPI	−2.216	−3.159	−3.874‡	2.441	1.164	0.805
United Kingdom						
RPIX	−1.279	−1.104	−1.499	1.796	1.671	0.673
RPI	−1.391	−1.509	−2.314	1.224	0.889	−0.332

Notes: Critical values were computed using the response surfaces in MacKinnon (1991). The superscripts ‡ and † denote the rejection of the hypothesis of a unit root at the 5 and 10 percent levels, respectively. For the United Kingdom, the RPIX and RPI are spliced with the appropriately rescaled CPI from 2004:M1 onwards.

process and the alternative is a covariance-stationary autoregression with neither a constant nor a trend. Notice that the null hypotheses are what inflation targeting predicts—and the alternatives are what price-level targeting predicts—about the price level. Since the ADF is a Wald-type test, the estimated processes are the ones specified under the alternative. The benchmark level of augmentation was selected using recursive t-tests, but the table reports results using other levels of augmentation to evaluate the robustness of results.

Table 2 reports results of KPSS stationarity tests. The KPSS test interchanges the null and alternative hypotheses in unit-root tests. Thus, the null hypotheses are what price-level targeting predicts—and the alternatives are what inflation targeting predicts—about the price level. Since the KPSS is a Lagrange multiplier test, the estimated processes are the ones specified under the null. The estimator of the residual spectrum at frequency zero uses a Bartlett kernel with Andrews (1991) data-based automatic bandwidth, but the table reports results using quadratic and Panzer kernels to evaluate the robustness of results.

In the case of Canada, the hypothesis that the CPI has a unit root can be rejected in the benchmark case, but this conclusion is not robust to using other levels of augmentation in the ADF test. On the other hand, the hypothesis that the CPI is covariance stationary cannot be rejected using the KPSS test regardless of the kernel used. Thus, these two tests lead to different conclusions regarding the presence of a unit root in the Canadian price level.

The ADF test is more powerful when applied to the CPI deviations from the implied path because the restrictions that this series has mean zero and no trend are imposed. It is important to recall that the trend used to construct the deviations is not estimated, but it is the one implied by the inflation target (that is, the CPI in January 1996 growing at the rate of 2 percent per year). In this case, the hypothesis that the deviations have a unit root can be rejected for all levels of augmentation, while the hypothesis that it is covariance stationary cannot be rejected for any kernel.[8]

[8]Indeed, in unreported work, I ran ADF tests using lag lengths from one to twelve lags and found that, except in the case where its number is nine, one can reject the null hypothesis of a unit root at standard significance levels.

Table 2. KPSS Stationarity Tests

	Variable					
	Price Level			Deviation from Implied Path		
		Robustness to Kernel			Robustness to Kernel	
Country and Price Index	Bartlett	Quadratic	Panzer	Bartlett	Quadratic	Panzer
Australia						
All Groups	0.133†	0.144†	0.180‡	0.341†	0.473‡	0.551‡
Canada						
CPI	0.121	0.101	0.105	0.213	0.173	0.174
Core CPI	0.097	0.082	0.082	0.461†	0.358†	0.376†
New Zealand						
All Groups	0.139†	0.175‡	0.219‡	0.616‡	1.621‡	1.805‡
Sweden						
CPI	0.134†	0.136†	0.142†	0.391†	0.508‡	0.590‡
United Kingdom						
RPIX	0.110	0.093	0.111	1.688‡	19.51‡	19.76‡
RPI	0.100	0.084	0.103	1.486‡	13.43‡	13.76‡

Notes: Critical values were taken from Kwiatkowski et al (1992). The superscripts ‡ and † denote the rejection of the hypothesis of covariance stationarity at the 5 and 10 percent levels, respectively. For the United Kingdom, the RPIX and RPI are spliced with the appropriately rescaled CPI from 2004:M1 onwards.

In summary, while acknowledging the inherently difficult task of making statistical statements about the long-run properties of a series on the basis of a relatively short sample, these results cast doubt on the view that the Canadian price level has a unit root as predicted by bygones-are-bygones inflation targeting.[9] Instead, it appears that the actual application of inflation targeting in Canada may induce statistical properties in the price level similar to those implied by a price-level targeting regime. I stress that these results do not mean that Canada covertly pursues a price-level targeting policy. In principle, it is possible that the observational equivalence between inflation and price-level targeting is the result of symmetric shocks ("luck"). That is, the central bank follows a bygones-are-bygones inflation-targeting policy, but symmetric shocks take inflation sometimes above, sometimes below, the target. However, the results reported above do not support this interpretation of the data because in this scenario the price level and its deviation from the implied path would still have a unit root (see, also, Kamenik et al. 2013). Instead, a more likely explanation is that the application of the inflation-targeting policy itself accounts for the results reported here. Possible mechanisms are those suggested by Nessén and Vestin (2005), Gorodnichenko and Shapiro (2007), Murchison (2010), and Kamenik et al. (2013). Nessén and Vestin show that a policy that targets the average level of inflation delivers time-series implications in between period-per-period inflation targeting and price-level targeting.[10] Gorodnichenko and Shapiro argue that a commitment to reverse the effect of policy mistakes on inflation

[9]Tables 1 and 2 also show that results are not robust to using the Core CPI, whose inflation rate is not targeted by the central bank. However, price-level targeting only generates predictions concerning the targeted index and is silent about indices that have different trends as a result of persistent movements in relative prices. In unreported work, I constructed a non-Core CPI using the relation CPI = (non-Core CPI)$^{0.17}$ (Core CPI)$^{(1-0.17)}$, where 0.17 is the sum of the weights of the eight product categories excluded from the CPI and which are reported in footnote 7. The hypothesis that non-Core CPI deviations from the implied path have a unit root cannot be rejected, but I found evidence that Core and non-Core deviations are cointegrated. A more detailed analysis of trends in the Canadian CPI components is left for future research.

[10]However, it would be empirically difficult to distinguish between period-per-period, average, and price-level targeting using their model because the number of periods over which the average is taken is a free parameter.

would empirically resemble price-level targeting. Using U.S. data, they estimate a model where the price level is one of the determinants of the interest rate in a Taylor rule and report that its coefficient is positive and statistically significant for the period when Alan Greenspan was Chairman of the Federal Reserve. It is interesting to note that Kamenik et al. (2013) find a similar result using Canadian data. Thus, it is possible that an error-correction mechanism akin to that described by Gorodnichenko and Shapiro may account for the "as if" price-level targeting regime in Canada. Finally, Murchison (2010) argues that interest rate smoothing induces history dependence in an inflation-targeting regime and, thus, may also explain the predictability of the Canadian price level.

In the cases of Australia, New Zealand, Sweden, and the United Kingdom, one cannot reject the hypotheses that the price level and its deviation from the implied path have a unit root,[11] but one can reject the hypothesis that these variables are covariance stationary. In the case of the United Kingdom, the result is driven by the sharp departure of the price level from the implied path after the beginning of the crisis. Overall, the statistical properties of the price levels in these four countries appear consistent with the view that inflation targeting allows shocks to have permanent effects on the price level.

2.4 Forecasting

This section asks whether forecasts of the price level are more accurate under either of two alternative models. In the first model, inflation targeting accommodates price-level shocks and bygones are bygones. Hence, starting from the current value, the future price level is expected to grow at the targeted inflation rate (e.g., as in panel B of figure 1):

$$E_t(\log(P_{t+h})) = (h/12)\log(1+\pi^*/100)\log(P_t), \tag{2}$$

where E_t is the expectation conditional on information available at time t, P_t is the price level at time t, and h is the horizon. In the

[11] It may be argued that results for Australia and New Zealand reflect low test power because their data is only available at the quarterly frequency and, hence, the number of observations is smaller than that of the other countries. However, Shiller and Perron (1985) show that the power of unit-root tests depends on the span of the data and not on the frequency at which the data is sampled.

second model, inflation targeting actually induces price-level stationarity around the path implied by the target. Thus, starting from the current value, the future price level is expected to eventually revert to this trend (e.g., as in panel F of figure 1):

$$E_t(\log(P_{t+h})) = ((t+h)/12)\log(1+\pi^*/100)\log(P_s), \qquad (3)$$

where P_s is the price level at time s, when the inflation-targeting policy began.[12] Table 3 reports the root mean squared error (RMSE) of price-level forecasts under both models at horizons of three, twelve, twenty-four, and thirty-six months. Note that by construction, forecasts are out of sample.[13]

For Canada, model 1 (bygones are bygones) delivers the smallest RMSE for CPI forecasts at horizons of a year or less, while for horizons longer than a year, it is model 2 (reversion to trend) that delivers the smallest RMSE. Thus, an agent forecasting the Canadian CPI more than a year ahead would do better by setting her forecast equal to the value implied by a deterministic path that grows at 2 percent per year starting from the CPI in January 1996 than by one starting from the CPI in the current period. In other words, the agent would produce a more accurate forecast if she were to treat the data as coming from a price-level targeting regime than from a bygones-are-bygones inflation-targeting regime. However, the reverse is true for shorter horizons. A possible explanation for this result is that the mean lag of the price-level deviations from the implied path is fourteen months and, thus, in the short term the CPI may still be far away from the implied path.

Results for Australia are similar, although the trend-reverting forecast becomes marginally superior at horizons longer than three years. In contrast, for New Zealand, Sweden, and the United Kingdom, the bygones-are-bygones model delivers the smallest RMSE at all horizons considered. This result is due to the fact that, as we saw in figure 2, the price level in these countries does not follow the trend implied by the inflation target.

[12] In the cases of Australia and New Zealand, I substitute $t/12$ with $t/4$ in (2) and $(\tau+t)/12$ with $(\tau+t)/4$ in (3).

[13] Using the mean absolute error (MAE) leads to the same conclusions and is not reported here to save space.

Table 3. RMSE of Price-Level Forecasts at Different Horizons

	Model 1				Model 2			
	Bygones are Bygones				Reversion to Trend			
Country and Price Index	**3**	**12**	**24**	**36**	**3**	**12**	**24**	**36**
Australia								
All Groups	0.565	1.339	1.994	2.381	2.240	2.280	2.332	2.379
Canada								
CPI	0.676	0.916	1.133	1.185	0.918	0.938	0.961	0.933
Core CPI	0.380	0.485	0.747	0.972	2.201	2.251	2.323	2.395
New Zealand								
All Groups	0.508	1.277	2.001	2.671	7.371	7.500	7.682	7.879
Sweden								
CPI	0.748	1.409	2.208	2.885	8.449	8.630	8.885	9.149
United Kingdom								
RPIX	0.602	0.881	1.493	1.981	2.350	2.392	2.454	2.520
RPI	0.638	0.999	1.569	2.045	2.142	2.169	2.204	2.258

Notes: Model 1 is $E_t(\log(P_{t+h})) = (h/12)\log(1+\pi^*/100)\log(P_t)$ and model 2 is $E_t(\log(P_{t+h})) = ((t+h)/12)\log(1+\pi^*/100)\log(P_s)$ with all notation as defined in the text. The RMSE is multiplied by 100 to reduce the number of digits to be reported. The horizon is in months. For the United Kingdom, the RPIX and RPI are spliced with the appropriately rescaled CPI from 2004:M1 onwards.

2.5 Expectations

The graphical and statistical evidence reported above suggests that the Canadian inflation-targeting regime may induce covariance stationarity and predictability in the price level. An important issue, however, is whether economic agents incorporate this empirical observation into their expectations and, perhaps, treat the price level as stationary around a time trend. This matters because the welfare benefits of price-level targeting arise from stabilizing price-level expectations. In order to examine this issue, I use data from two surveys on inflation expectations.

The first survey is the Canadian Survey of Economic Expectations (CSEE) carried out by Towers Watson. Every November, this survey asks a sample of portfolio managers and business economists their forecasts of the rate of CPI inflation in the incoming calendar year. Since Statistics Canada releases the CPI for the previous month in the second half of each month, I assume that when the survey is carried out, respondents know the September, but not yet the October, CPI.[14] With this assumption and the reported CPI inflation expectations, it is possible to construct an implicit expectation for the price level itself. For this project, I use the ten surveys starting with the one for 2002. (Unfortunately, surveys prior to 2002 do not appear to be available.)[15]

For each year in the sample, figure 3 plots a fan chart with CPI expectations (thin lines), the realized price level (thick line), and the price-level path implied by the inflation target (dashed-dotted line). For example, the panel for 2002 plots the expectations in November 2001 about the CPI in December 2002 with the CPI of September 2001 as the starting point. The five thin lines correspond to the tenth, twenty-fifth, fiftieth, seventy-fifth, and ninetieth percentiles of the forecasts.[16] The raw data are reported in table 4.

[14]For the two years when the survey was carried out in December (2004 and 2005), I assume that respondents know the October, but not yet the November, CPI. Also, there are two years (2007 and 2011) for which the documentation does not indicate the month when the survey was carried out. For the purpose of this analysis, I will assume that it was carried in November, but assuming December instead does not change the conclusions.

[15]Recent surveys are freely available at www.towerswatson.com/en-ca.

[16]For the years up to 2005, the survey reports the lowest and highest forecast, rather than the tenth and ninetieth percentiles.

Figure 3. Price-Level Expectations in Canada

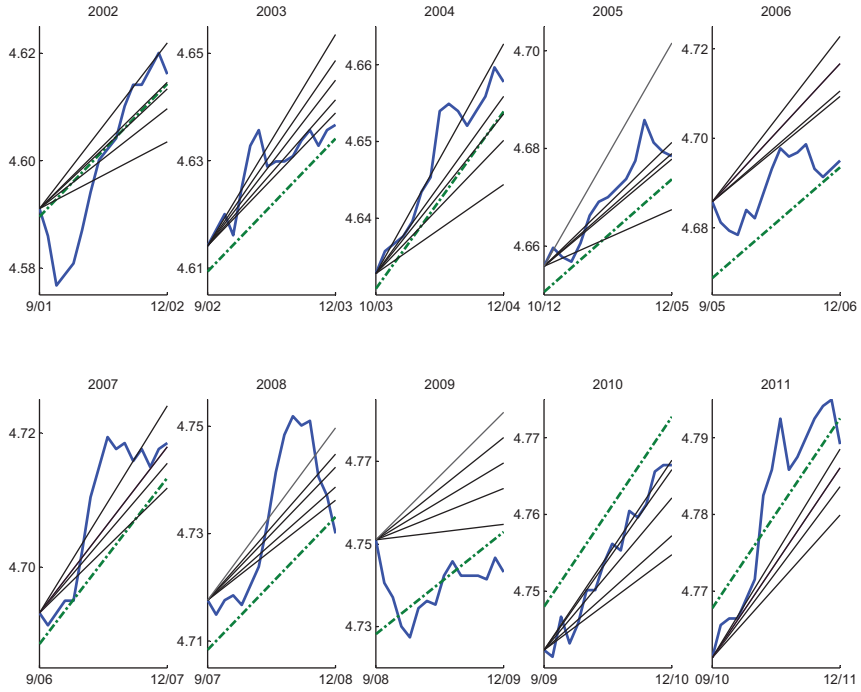

Notice that the expectation that the CPI will continue to grow at the (average) rate of 2 percent per year for the next fifteen months is contained in the fan chart for all years and coincides with the median forecast in four out of the ten years. This is what one would expect to see under the bygones-are-bygones version of inflation targeting. On the other hand, there is a considerable proportion of respondents that expect a return to trend in the near future, and their forecasts tend to be generally closer to the actual outcomes in these plots. In the surveys from 2002 to 2009, when the price level is above the implied path at the time the survey is carried out, the proportion of respondents that anticipate an inflation rate below 2 percent and, thus, a reversion to the path, is substantial. For example, when the 2008 survey was carried out, the price level was 0.92 percent above trend and all respondents below the 50th percentile anticipated inflation rates below 2 percent (see table 4), meaning that they thought that the price level would approach the trend implied by the inflation

Table 4. Inflation Expectations in Canada

Year	Number of Respondents	Deviation (in %)	Percentile				
			10	25	50	75	90
2002	41	0.15	1.0	1.5	1.8	1.9	2.5
2003	37	0.48	2.0	2.2	2.5	2.8	3.2
2004	54	0.20	1.0	1.5	1.8	2.0	2.6
2005	40	0.53	1.0	1.9	2.0	2.2	4.0
2006	46	1.71	1.9	2.0	2.5	2.5	3.0
2007	38	0.46	1.5	1.8	2.0	2.0	2.5
2008	32	0.92	1.5	1.7	2.0	2.2	2.6
2009	40	2.28	0.3	1.0	1.5	2.0	2.5
2010	44	−0.57	1.0	1.2	1.6	1.9	2.0
2011	43	−0.65	1.5	1.8	2.0	2.0	2.2

Note: The raw data were taken from various issues of the Canadian Survey of Economic Expectations carried out by Towers Watson.

target during the forecast horizon. The 2003 survey is an exception to this general observation because even forecasters in the lowest tenth percentile anticipate an inflation rate of 2 percent and the rest anticipate inflation rates larger than 2 percent.

A drawback of the CSEE is that the survey is taken only once a year. For this reason, I also use a second survey, namely the Consensus Forecasts (CF), produced by Consensus Economics. In the first few days of each month, this survey asks professional economists their inflation forecasts for the current and next calendar year. Hence, this survey provides forecasts of the same variables at different horizons and with different information sets. Unfortunately, these data are proprietary and I was only able to use (with permission) the mean forecast in each month between 2002:M1 and 2013:M3.[17] As before, I assume that when the CF survey is carried out, respondents know the CPI of the month before the previous one.

[17]There is some overlap in the institutions included in the CSEE and CF, but there is no way to ascertain if the same individual answers both surveys. Also, note that the sample of institutions in CF is smaller (between fifteen and twenty institutions) than the sample in the CSEE (between thirty and forty institutions).

Figure 4. Inflation Expectations in Canada

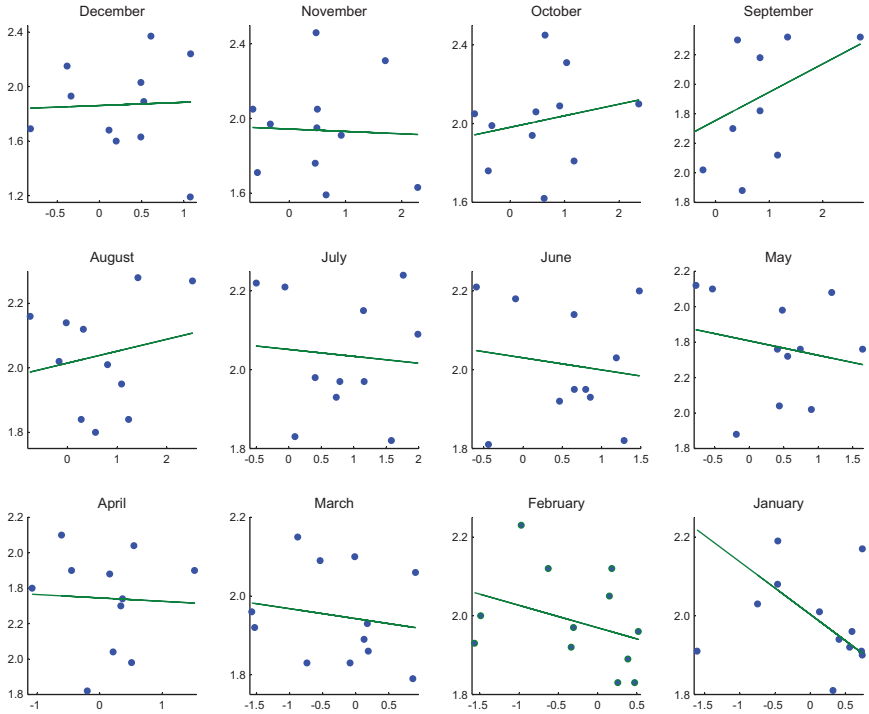

Figure 4 plots forecasts of inflation in the next calendar year as a function of the price-level deviation from the target when the survey was taken and controlling for the forecast horizon. Thus, the figure contains twelve panels—one for each month of the year—and the points in each panel (say, May) are the (mean) forecasts of surveys taken that month in all years (say, May 2002, May 2003, etc.). Thus, the forecast horizon is the same in each panel, and as we move in the figure from December to November to October and so on, the forecast horizon increases. The price-level deviation is that known to respondents when the survey was taken. Finally, the thick line is the fitted value of a robust linear projection of the forecasts on the price-level deviations and a constant. A negative slope means that when the price level is below (above) trend, agents forecast relatively high (low) inflation and, consequently, the price level is expected to come back to trend. The converse holds when the slope is positive.

It is interesting to note that as the forecast horizon increases, the slope decreases and becomes negative for July forecasts and earlier. Thus, roughly speaking, when market participants are surveyed in January–July to forecast inflation in the next calendar year (that is, to forecast the price level twenty-four to eighteen months ahead), their forecasts suggest an expectation of trend reversion in the price level. Instead, when they are asked in August–December (and, thus, the horizon for their price-level forecast is seventeen to thirteen months ahead), their forecasts do not suggest an expectation of trend reversion. These results are in line with results in the previous section that showed that for longer forecast horizons, a trend-reversion forecasting model has smaller RMSE than a bygones-are-bygones model, with the opposite being true for shorter horizons.

However, although these results are suggestive, note that in none of these panels is the slope statistically different from zero, but in all panels the intercept is not statistically different from 2 percent. In other words, one would not reject the hypothesis that regardless of the price-level deviations from the implied target, market participants forecast an inflation rate equal to the announced target, as predicted by bygones-are-bygones inflation targeting.[18]

In summary, the survey data only provide suggestive evidence that agents treat the Canadian price level as a trend-reverting process. Instead, as a whole, these data are consistent with the view that agents forecast the price level as a bygones-are-bygones process. Thus, it is possible that an explicit price-level targeting policy may be welfare improving if it helps stabilize price-level expectations around a pre-announced path.

3. Discussion

This paper reports evidence that in Australia, New Zealand, Sweden, and the United Kingdom, the price level drifts away from the path implied by the inflation target and is well described by a unit-root

[18]It is important to stress that long-term inflation forecasts (available in both surveys) are not informative as to whether agents view the inflation-targeting policy as bygones are bygones or as an "as if" price-level targeting regime because both policies, if credible, predict that the growth rate of the price level will be 2 percent.

process. This evidence is consistent with the view that, under inflation targeting, shocks have permanent effects on the price level. This paper also reports evidence that in Canada, the price level closely tracks the path implied by the inflation target and has statistical properties similar to those implied by a price-level targeting regime.

This evidence does not mean that Canada covertly pursues price-level targeting. Instead, the evidence means that there is heterogeneity in the actual application of inflation targeting across countries, and that the description of inflation targeting as a policy that systematically accommodates price-level shocks may be restrictive. In fact, the Canadian experience suggests that inflation targeting may be flexible enough to deliver a stationary and predictable price level.

In interpreting these results, it is important to keep in mind two caveats. First, since it is difficult to make statistical statements about the long-run properties of a series on the basis of a relatively short sample, results from unit-root tests are best taken as suggestive. Second, the experience of the United Kingdom illustrates the case of an inflation-targeting regime that up to the financial crisis featured a predictable price level but that (so far) appears to have accommodated the effects of that large shock. This indicates that the size of shocks matter and that an inflation-targeting central bank may choose to accommodate small, but not large, shocks.

References

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

Andrews, D. W. K. 1991. "Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation." *Econometrica* 59 (3): 817–58.

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

Bernanke, B., T. Laubach, F. Mishkin, and A. Posen. 1999. *Inflation Targeting*. Princeton, NJ: Princeton University Press.

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

Dodge, D. 2005. "Our Approach to Monetary Policy: Inflation Targeting." Remarks to the Regina Chamber of Commerce, Regina, Saskatchewan, Canada, December 12. Available at http://www.bankofcanada.ca/2005/12/publications/speeches/approach-monetary-policy-inflation-targeting/.

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

Gorodnichenko, Y., and M. D. Shapiro. 2007. "Monetary Policy when Potential Output Is Uncertain: Understanding the Growth Gamble of the 1990s." *Journal of Monetary Economics* 54 (4): 1132–62.

Kamenik, O., H. Kiem, V. Klyuev, and D. Laxton. 2013. "Why Is Canada's Price Level So Predictable?" *Journal of Money, Credit and Banking* 45 (1): 71–86.

Kwiatkowski, D., P. C. B. Phillips, P. Schmidt, and Y. Shin. 1992. "Testing the Null Hypothesis of Stationarity against the Alternative of a Unit Root." *Journal of Econometrics* 54 (1–3): 159–78.

MacKinnon, J. 1991. "Critical Values for Cointegration Tests." In *Long-Run Economic Relations: Readings in Cointegration*, ed. R. F. Engle and C. W. J. Granger, 267–76. Oxford: Oxford University Press.

Melino, A. 2012. "Inflation Targeting: A Canadian Perspective." *International Journal of Central Banking* 8 (S1): 105–31.

Murchison, S. 2010. "Price-Level Targeting and Relative-Price Shocks." *Bank of Canada Review* (Summer 2010): 11–21.

Nessén, M., and D. Vestin. 2005. "Average Inflation Targeting." *Journal of Money, Credit and Banking* 37 (5): 837–64.

Parkin, M. 2009. "What Is the Ideal Monetary Policy Regime?" C. D. Howe Institute Commentary No. 279.

Shiller, R. J., and P. Perron. 1985. "Testing the Random Walk Hypothesis: Power versus Frequency of Observation." *Economics Letters* 18 (4): 381–86.

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

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

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

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

Woodford, M. 1999. "Optimal Monetary Policy Inertia." NBER Working Paper No. 7161.