Two Types of (Slight) Flexibility in Bank of Canada Projections, 2003–2019

Gregor W. Smith[†]

May 2019

Abstract

Central bankers and other economists sometimes use the word *flexible* to describe the inflation targeting practiced in Canada. The flexibility refers to deferring the planned return of the CPI inflation rate to the 2% target. This paper outlines two ways to measure flexibility, by seeing whether the forecast for future inflation varies with either (1) current values of inflation or output growth or (2) the forecast for future output growth. The latter correlation would suggest that the Bank of Canada aims to stabilize output growth in addition to stabilizing the inflation rate. I describe how to detect evidence of these two types of flexibility in the Bank of Canada's own forecasts (also known as projections) from its quarterly *Monetary Policy Report*. But there is little sign of economically significant flexibility between 2003 and 2019.

JEL classification: E58.

Keywords: monetary policy, forecasts, inflation targeting

[†]Department of Economics, Queen's University, Canada; smithgw@econ.queensu.ca. I thank the Social Sciences and Humanities Research Council of Canada for support of this research. For helpful comments I thank Jeremy Kronick, David Longworth, John Murray, and a referee of this journal.

1. Introduction

Since the 1960s, economists' views of how monetary policy works have focused on the role of expectations. For example, there is evidence that expectations about future inflation influence the current inflation rate. Thus central banks operate in part by managing expectations, and so their own forecasts are one of their policy tools. Forward guidance about short-term interest rates, a strategy used by the Bank of Canada during the financial crisis, thus is in the toolkit.

Since 2003, the Bank of Canada has published its projections for CPI inflation and real output growth in its quarterly *Monetary Policy Report* (*MPR)*. The *MPR* typically shows the inflation rate converging to 2%, the mid-point of the Bank's target range. At the same time, the Bank describes its inflation targeting as *flexible*, meaning that the return to 2% may sometimes be deferred. Flexible inflation targeting matters to practitioners because it affects conditional forecasts for inflation and policy interest rates. One would like to predict these in advance and not necessarily wait for the central bank's predictions *i.e.* to predict when the Bank will be flexible.

I argue that the Bank's projections implicitly provide information about the degree of flexibility in monetary policy. I describe two forms of flexibility, or ways to detect it. First, the projected value for inflation at a given horizon in the future may vary with some current, public data, such as the current inflation rate. Second, the projected value may vary with another of the Bank's projections, and particularly the projected growth rate of real output. Using the projections in the *Monetary Policy Report*, I provide simple statistical tests for each type of flexibility. I find that flexibility in inflation targeting can be detected, but it is on a very small scale. Historically, then, I suggest that the Bank of Canada can be described simply as having targeted the inflation rate.

This study is not about the accuracy of the projections. Realized inflation and output growth of course would be needed to assess the Bank's track record, but they also are affected by a range of shocks, so their values do not always shed light on the Bank's goals. This study also is not about the desirability or efficacy of flexibility in inflation targeting, though I cite macroeconomic studies on that topic. It also is not about the statements of Bank of Canada officers, though those may play a role in influencing expectations also.

However, I do next give some examples of those statements, so the reader may see that references to flexibility are prominent there.

2. Two Types of Flexibility

The preamble to the October 2018 *Monetary Policy Report* says (with italics in the original):

> Canada's inflation-targeting framework is *flexible*. Typically, the Bank seeks to return inflation to target over a horizon of six to eight quarters. However, the most appropriate horizon for returning inflation to target will vary depending on the nature and persistence of the shocks buffeting the economy.

Bank of Canada officials have described three aspects of flexibility. First, it involves a deferral in the planned return of the inflation rate to 2%, the mid-point of the target range *i.e.* a longer horizon before 2% is reached. Second, it is applied in response to the need to use monetary policy to address some other goal, such as financial stability or unemployment. For example, if inflation were below target but the Bank viewed financial leverage as excessive, it might keep its policy interest rate higher than it otherwise would, and so delay an increase in the inflation rate. Likewise if inflation were above target but the Bank viewed the unemployment rate as abnormally high, then it might delay raising that interest rate. Third, the Bank has been well aware of the possible tradeoff between flexibility and credibility, and hence of the need to clearly communicate deferrals. Koeppl (2009) also discusses how the exercise of flexibility requires extra communication to preserve credibility.

The Bank of Canada (2011, 2016) has described this perspective in detail in the background documents at renewals of the inflation target. For example, the 2016 documents describe lengthening the horizon in April 2009. Its senior officials also have described flexible inflation targeting regularly in speeches, for example those by Carney (2008, 2009, 2012a, 2012b, 2013), Macklem (2014), and Poloz (2013).

While I read these documents and speeches with interest, I here study only the data on the Bank's projections. One might think it would make sense to judge deferrals simply by recording the first horizon at which the forecasts return to 2% and then seeing if that horizon varied over time. In practice though, one might not want to count a forecast of

1.8% or 1.9% inflation as a deferral. Thus I did not record that horizon.

Instead, note that a deferral in returning to 2% naturally implies that the projected inflation rate will differ from 2% at some more standard, targeted horizon (perhaps 6–8 quarters, as noted in the quotation above). The first type of flexibility—labelled FIT 1.0—involves this departure from 2% covarying with some current information. I use the current and lagged projections for inflation and output growth as examples of that current information, and test for this covariation.

Label the projection at quarter t for inflation h quarters later as $P_t\pi_{t+h}$. Then the first tests use these linear statistical models, with lagged inflation:

$$P_t\pi_{t+h} = \beta_0 + \beta_1\pi_{t-1} + \epsilon_{t,h}, \tag{1a}$$

and then with current inflation:

$$P_t\pi_{t+h} = \beta_0 + \beta_1\pi_t + \epsilon_{t,h}. \tag{1b}$$

The parameter β_1 measures the degree to which the projection, at horizon h, varies with lagged or current inflation. I also test whether the forecasts also vary with the nowcast of lagged output growth, y_{t-1}:

$$P_t\pi_{t+h} = \beta_0 + \beta_1\pi_{t-1} + \beta_2 y_{t-1} + \epsilon_{t,h}, \tag{2a}$$

or with the nowcast of current output growth y_t:

$$P_t\pi_{t+h} = \beta_0 + \beta_1\pi_t + \beta_2 y_t + \epsilon_{t,h}. \tag{2b}$$

The parameter β_2 measures whether the inflation projection depends on lagged or current output growth. In each of these statistical models the 2% target is subsumed in the intercept so there is no need to define the departure from it on the left-hand side.

The second usage—here labelled FIT 2.0—formally defines flexibility as targeting a second goal in addition to the inflation rate. As described by macroeconomists, this second goal usually involves stabilizing projected real output growth. Of course the output projection itself depends on current information, but that is not observed outside the Bank of Canada. The projection for output growth, denoted $P_t y_{t+h}$, pools a large amount of information that is not available to outside observers.

Here the central bank has a loss function that weights both departures from its inflation target π^* and departures of output growth from a target y^*. Minimizing that loss function at horizon h gives rise to this condition:

$$P_t \pi_{t+h} - \pi^* + \phi(P_t y_{t+h} - y^*) = 0 \tag{3}$$

or, rearranging:

$$P_t \pi_{t+h} = \pi^* + \phi y^* - \phi P_t y_{t+h}. \tag{4}$$

Sometimes the output-stabilization goal instead is described as minimizing an 'output gap' betwen the level of GDP and a long-run trend. But the *MPR* does not typically contain estimates of such a gap or numerical forecasts of its future values, so I focus on the growth rate.

At short horizons or nowcasts one might expect the value of ϕ to be negative. In other words, there will be a positive relationship between inflation and output growth, sometimes described as a Phillips curve. But at longer horizons, which are targeted by monetary policy, we expect ϕ to be positive. The logic is straightforward. We imagine that tightening monetary policy by increasing the overnight interest rate tends to lead to lower inflation and also lower output growth at some horizon h. If the projection involves output growth that is unusually low then the central bank will slightly compromise its inflation goal in order to stabilize output growth. In other words it will be less inclined to raise the short-term interest rate it controls and so will risk inflation's being above target. As a result, the forecast for inflation will be higher than it otherwise would be.

This strategy is completely symmetric. Conversely, then, if the projection involves an unusually high value for output growth then the central bank will again be concerned to stabilize that growth rate. As a result it will be more inclined to raise the short-term interest rate it controls, even if inflation is at or near the 2% target. It thus will risk the inflation rate's being below the 2% target.

This characterization of inflation targeting is called a *targeting rule*, as opposed to an *instrument rule* which describes the reaction function for a short-term interest rate.

Articles by Svensson (2003) and Woodford (2004, 2007a, 2007b) are classic references on the formulation (and desirability) of targeting rules. Svensson (2009, p 2) describes flexible inflation targeting like this:

> In this framework, the central bank takes financial conditions such as credit growth, asset prices, imbalances, potential bubbles and so on into account only to the extent that they have an impact on the forecast of inflation and resource utilization. Inflation and resource utilization are target variables here, that is, variables that are arguments of the central bank's loss function.
>
> Hence, financial conditions are not target variables. Instead, they are only indicators, as they provide information about the state of the economy, the transmission mechanism and exogenous shocks to the central bank. Financial conditions then affect policy rates only to the extent that they have an impact on the forecast of inflation and resource utilization.

Kuttner (2003) first suggested using a central bank's own forecasts to study targeting rules. As he noted (p 95), this approach has three appealing features. First, it avoids the challenge of forming one's own forecasts of inflation and output growth, measures which can be very difficult to predict at longer horizons. Second, central bank forecasts may be better than econometric forecasts. They can embody a range of real-time indicators. And in some cases, such as the Greenbook forecasts of the Federal Reserve Board, their relative accuracy has been documented. Third, they presumably embody a central bank's realistic interest-rate reactions, even though the Bank of Canada does not publish a projected path for the overnight interest rate.

Kuttner reported correlations between forecasts of inflation and an output gap, for the Bank of England, Reserve Bank of New Zealand, and Riksbank. Otto and Voss (2011) studied forecasts by the Reserve Bank of Australia, which has an explicit, dual mandate to stabilize inflation and output. They found a large and significant $\hat{\phi}$, especially when they allowed the target π^* to vary within the target range. But, as a pioneering study, they had forecast data for only 13 quarters and 5 horizons.

Some studies instead begin with realized outcomes for π_{t+h} and y_{t+h} and then construct their own forecasts using variables known at time t. Rowe and Yetman (2002) developed this method and applied it to Canada for the period from the mid 1970s to the late 1990s. They found that departures of inflation from the 2% midpoint of the target range could not be predicted, after the advent of inflation targeting. Otto and Voss (2014) followed this method to estimate ϕ for Canada over the 1996–2007 period. They also cited studies for the US that use realized measures to assess the goals of the Federal Reserve. They found a positive and significant $\hat{\phi}$. Their results are not directly comparable to those based on Bank of Canada projections, though, because they used core inflation and a quarter-to-quarter growth rate for output.

In the projections, the statistical test for FIT 2.0 then takes place in this regression:

$$P_t\pi_{t+h} = \alpha - \phi P_t y_{t+h} + \epsilon_{t,h} \tag{5}$$

which I estimate by ordinary least squares. The error term $\epsilon_{t,h}$ reflects other events that may affect the inflation forecast but are unrelated to the output growth forecast. These terms may be correlated across time and horizons so standard errors for the coefficients are estimated with a Newey-West estimator that is heteroskedasticity-and-autocorrelation-consistent (HAC).

Notice that the original target-tradeoff (4) implies that in the regression model (5) $\alpha = \pi^* + \phi y^*$. If we assume that the target for inflation in Canada is $\pi^* = 2$ then we can invert the estimated intercept to find the implicit target for output growth: $\hat{y}^* = (\hat{\alpha} - 2)/\hat{\phi}$.

In sum, I use the Bank of Canada's definition of flexibility, defined as allowing for deferrals in the return of projected inflation to 2%. Such a deferral implies that at the usual horizons for returning inflation to target (such as 6–8 quarters) $P_t\pi_{t+h}$ will not equal 2. Then I look for evidence of these departures by seeing whether $P_t\pi_{t+h} - 2$ varies with current conditions (FIT 1.0) or with output growth projections (FIT 2.0).

3. *MPR* Projections

To measure flexibility, I first collected projections from the Bank of Canada's *Monetary Policy Report*. The *MPR* was first issued in 1995, and appeared twice each year until 1999. From 1999 to 2009 it was accompanied by an *Update* in January and July of each year. Since July 2009 the quarterly issues have all been called the *MPR*. They appear in January, April, July, and October of each year, thus in the first month of each quarter.

The projections appear in a table entitled 'Projection for Core and Total CPI Inflation' (from April 2003 to April 2005), then 'Projection Summary' (for July 2005), then 'Summary of the base-case projection' (October 2005 to April 2013), and later 'Summary of the projection for Canada' (since July 2013). I use the terms 'projection' and 'forecast' interchangeably. I focus on the Bank of Canada's target, which is the year-over-year percent change in the total CPI, denoted π_t. Its forecasts first appeared in the April 2003 issue. These data span 66 reports for π_t up to April 2019.

I also collected forecasts for the year-over-year percent change in real GDP (denoted y_t) which first appeared in July 2005. These data span 57 reports up to April 2019. Binette and Tchebotarev (2017) study the accuracy of *MPR* forecasts for real output growth for a longer time span (1997–2017) by focusing on only annual forecasts from the April and October reports and also reading forecasts prior to 2005 from the text of the report. I focus on a shorter time span where one can match up the quarterly forecasts for both inflation and real GDP growth.

Forecasts apply to the just completed quarter, the current quarter, and then a range of future time periods. Until April 2009 the future periods included several quarters followed by half years and full calendar years. In those cases I attribute the forecasted growth rate to each relevant quarter. Since April 2009 the projections have been only for quarters. The number of quarters forecasted (or greatest horizon) varies over time because of calendar year coverage. For example, the April *MPR* begins to cover the entire subsequent calendar year and so often has contained more forecasts than the reports in October or January.

I use the term 'current quarter' to refer to the quarter in which the *MPR* is issued: Q1 for January, Q2 for April, Q3 for July, and Q4 for October. Notice that the relevant quarter has just begun at the time the report is released. I use the term 'last quarter' to refer to the quarter before the one in which the *MPR* is released. Notice that those entries also can be nowcasts for CPI inflation, because in some cases the *MPR* publication date falls before the release date for the previous month's (and hence quarter's) CPI. For output growth, these are always nowcasts because the last quarter's GDP figure is not released until late in the second month or early in the third month of the current quarter.

There has been a recent change in the horizons in the *Monetary Policy Report.* From 2006 to 2015 the horizon varied from 8 to 11 quarters, depending on the quarter in which the Report was issued, because of calendar year coverage. Then in January 2016 the Report began to list forecasts only for the inflation rate and output growth rate in the fourth quarter of the year (It continues to list nowcasts for the previous and current quarter.) The January *MPR* gives this fourth-quarter projection for the current and next year, while the April, July, and October issues give them for the current year and the next two years. Thus fewer projections are tabulated than in earlier reports, and there are gaps in the sequence of horizons.

The Bank of Canada also makes available its staff projections with a five year lag. Champagne, Poulin-Bellisle, and Sekkel (2018) describe these data and use them to judge the relative accuracy of the forecasts. Similarly, the staff of the Federal Reserve Board of Governors produces forecasts for the US using assumptions about monetary policy, in the form of the Greenbook. These are presented at each meeting of the Federal Open Market Committee. They too are made public after a lag of five years. A number of studies (*e.g.* Faust and Wright, 2009) have examined their accuracy compared to other forecasts.

I do not use the Bank of Canada's staff projections to measure flexibility, for three reasons. First, unlike the projections in the *MPR* they are not the official view of the Bank's Governing Council. Second, there also is a timing difference, with the *MPR* projection based on an extra month of data. And Champagne, Poulin-Bellisle, and Sekkel (2018, section 5.2) show that the two sets of forecasts differ. Third, the staff projections were not (and are not) revealed at the time of monetary policy decisions. This paper focuses on flexibility as it could be communicated with any deferrals that occurred in *MPR*s.

Figure 1 graphs *MPR* projections for inflation, $P_t\pi_{t+h}$, since 2003, for each horizon h = 3–8. For $h = 3$ the largest departures from 2% are during 2008–2010, which is consistent with the renewal document's (Bank of Canada, 2016) mention of a deferral in 2009. Once $h = 4$ it is more difficult to argue that those observations are outliers, though. For $h = 5$ and $h = 6$ there are some values below 2% from 2003 to 2005. But thereafter the line appears steady. Certainly by $h = 6$ the 2008–2010 period is not an outlier.

You might think: The Bank of Canada won't forecast a return to 2% inflation at short horizons regardless of the initial inflation rate. And you would be correct. However, one would expect that, using a larger value of the horizon h, under *inflexible* inflation targeting $P_t\pi_{t+h} = 2$ regardless of the initial inflation rate. Flexible inflation targeting would simply make these somewhat longer-horizon diagrams look like short-horizon ones, in that larger departures from 2% are possible.

Overall, from figure 1 it is difficult to see much evidence of such a deferred return to 2% or, equivalently, departures from 2% at horizons of 6–8 quarters. For recent data, recall that realized inflation has generally been below 2%, due to a series of unfortunate events including declines in commodity prices. But the projections tend to be stable nevertheless. The next section formally tests for any pattern relating the inflation projections to lagged or current inflation and output growth at the time of the projection.

4. Statistical Tests for FIT 1.0

I first study the role of lagged or current inflation, as in equations (1a) and (1b). The results essentially summarize 16 scatterplots, one for each explanatory variable and horizon. Table 1 gives the coefficients $\hat{\beta}_0$ and $\hat{\beta}_1$. The table also reports the R^2 statistic and the number of observations N.

Brackets contain the standard errors for the coefficients. I note two technical details concerning these standard errors. First, they are HAC Newey-West with lag length equal to the forecast horizon minus one. That lag length is chosen because an h-step-ahead forecast induces a moving average of order $h-1$ in forecast errors. In practice, though, the standard errors are not sensitive to the lag length chosen. Second, graphing squared residuals vs. the right-hand-side-variable shows that they are largest when that x-variable is near its mean. In that case OLS standard errors are too large and robust ones are (correctly) smaller.

If a forecast differing from 2% reflected a projected gradual adjustment then $\hat{\beta}_1$ would be positive: a lagged or current inflation rate much above 2% would be associated with a projection that remained above 2%. There is indeed a positive association between the projection and lagged inflation at horizons 1–2 and between the projection and current inflation at horizons 1–3. As h increases, though, notice that the coefficient $\hat{\beta}_1$ falls and it becomes statistically indistinguishable from zero, while the intercept $\hat{\beta}_0$ rises and becomes indistinguishable from 2. Thus, if there is flexibility, the current conditions at work are not reflected in current or lagged inflation.

The dispersion in inflation projections at intermediate and longer horizons thus must reflect something else. Equations (2a) and (2b) include lagged or current output growth in addition to lagged or current inflation, as potential indicators. Here the sample begins in July 2005.

Table 2 contains the results. Now there are some statistically significant effects on the projections (as seen in the ratios of $\hat{\beta}_1$ and $\hat{\beta}_2$ to their standard errors or the larger R^2 statistics) at most horizons, including the longer ones. There is a significant role for lagged output growth or for current output growth. Moreover, controlling for this effect, there now is some evidence of a role for π_{t-1} even at $h = 7$ and $h = 8$. However, though these effects are statistically significant (for example at the 5% level of significance) they are very small. All the variables are in four-quarter growth rates, expressed in percentage points. Given the values of $\hat{\beta}_1$ and $\hat{\beta}_2$, then, the indicators have no economically significant effect on the inflation projections.

5. Statistical Tests for FIT 2.0

It also is possible that inflation projections are related to output growth projections at the *same* horizon, as suggested by FIT 2.0. Figure 2 shows output growth projections since 2005. These are more variable than the projections for CPI inflation and they show less tendency to stabilize as the horizon increases. For this second type of flexibility, table 3 provides estimates of the parameter ϕ, using data from each individual horizon 1–8 and then from various groups of horizons: 1–8, 4–8, 5–8, and 6–8. The table presents estimates $\hat{\phi}$ and their standard errors. It also presents estimates of the intercept $\hat{\alpha}$ and then unscrambles the implicit target for output growth $\hat{y}^*$ (assuming that $\pi^* = 2$) along with its standard error.

As section 2 noted, one expects $\hat{\phi}$ to be negative at short horizons, where inflation and output growth are positively correlated due to a Phillips-curve effect. When estimated with projections from an individual horizon, that is exactly what one finds for horizons 1–4, in the first four rows of table 3.

At longer horizons, which the Bank of Canada targets, we expect $\hat{\phi}$ to be positive, reflecting the tradeoff between its two stabilization goals. When estimated with projections from an individual horizon, $\hat{\phi}$ is positive for each $h = 5, 6, 7,$ or 8 and statistically different from zero at the 1% level in each case. When estimated with groups of horizons (to aid precision), $\hat{\phi}$ is positive and statistically significant at the 1% level for the groups of horizons 5–8 and 6–8.

But the influence of the output projection on the inflation projection is small. For example, consider $\hat{\phi} = 0.038$ (with standard error 0.008) estimated from horizon 8 in table 3. These values imply that a fall in the forecast of output growth by 1 percentage point for that horizon (a significant slowdown in economic growth) would lead to an upward shift in the inflation projection by 0.038 of a percentage point (with a 95% confidence interval of 0.022–0.054). This effect is detectable, but it is very small.

Notice also that the goodness-of-fit statistics R^2 are less than 0.50 in table 3. Inflation projections are not simply equal to a constant value of 2% (subsumed in the intercept α) and then adjusted for output-growth projections. So this characterization of the targeting rule does not provide a complete description of the projections.

Finally, one might wonder what happens if both current or lagged indicators and projections for future output growth are included as regressors. The results (not shown) vary with the horizon. At some horizons there is a statistically significant role for both types of indicators and at other horizons for neither. But $P_t y_{t+h}$ does not generally drive out the lagged or current indicators from the regression equation. However, once $h \geq 6$ all coefficients are less than 0.07. As a result, the effect of any combination of the two types of flexibility once again is very small.

6. Conclusion

It is possible there is some factor, missing from this study but known by the Bank at the time of projections, that has varied over time and explains the departures of inflation projections from the midpoint of the target range. Perhaps other researchers will update the data and investigate this possibility in the future.

But there are three challenges in identifying that factor. First, the pattern of departures of $P_t \pi_{t+h}$ from 2% in figure 1 suggests that they are not very persistent. Presumably, then, such a missing factor would have to have that time-series property also, as opposed to evolving slowly and persistently.

Second, this study found some statistically significant estimates, so it is not the case that there is simply not enough data to measure anything precisely. I did not file a pre-analysis plan, but I have reported the findings from each regression I ran, and I did not study any other covariates. However, with limited data it can be difficult to control test size in searching across indicators. Equivalently, there is a risk of false positives (or type I error), in other words a risk of falsely rejecting the null hypothesis that the parameter on an indicator is zero. This is simply a formal statement of the fact that one can perhaps search across many indicators and find one that happens to be randomly correlated with the projections. This is akin to the syndrome of 'explaining' fluctuations in the stock market or the exchange rate after the fact. This search for an explanatory variable is especially challenging with a small number of observations. Of course this data-mining issue is relevant to outside observers, not to the Bank of Canada. It knows what it did in the past and of course could communicate the exercise of flexibility in the future.

Third, the most plausible rationale for extra factors seems to be that they influence $P_t y_{t+h}$, as indeed recommended by Svensson (2009), as quoted above. But I find the variation in these output-growth forecasts to be insignificant economically in explaining departures from $P_t \pi_{t+h} = 2$. For example, if the Bank tracked some financial variables that it thought might destabilize output growth in the future, then presumably the fluctuations in those variables already would be reflected in the projections for output. I have tested the leading theory of flexibility and also prominent covariates from the projection summary table itself. Presumably if the Bank regularly conditioned on some other variable it would say so, in the interest of transparency and thus to enhance the effect of a deferral.

The idea behind flexible inflation targeting is that the deferral in returning to target yields some benefits today. It is interesting to think how one would evaluate the success of this strategy or assess its desirability. Figure 3 shows realized, four-quarter inflation (in black) and output growth (in grey) for Canada since 2003. Seeing the fluctuations in output—and especially the 2009–2009 recession—raises the question of whether monetary policy (perhaps in concert with fiscal policy) should have been more flexible. But, at least judging by these simple tests, flexible inflation targeting remains a largely untried policy in Canada.

References

Bank of Canada (2011) Renewal of the Inflation-Control Target: Background Information–November 2011. Ottawa: Bank of Canada.

___________(2016) Renewal of the Inflation-Control Target: Background Information–October 2016. Ottawa: Bank of Canada.

Binette, André and Dmitri Tchebotarev (2017) Evaluating real GDP growth forecasts in the Bank of Canada *Monetary Policy Report*. Bank of Canada Staff Analytical Note 2017-21.

Carney, Mark (2008) Flexibility versus credibility in inflation-targeting frameworks. 7th BIS annual conference, Lucerne Switzerland.

___________(2009) Some considerations on using monetary policy to stabilize economic activity. Federal Reserve Bank of Kansas City symposium, Jackson Hole, Wyoming.

___________(2012a) A monetary policy framework for all seasons. US monetary policy forum, New York.

___________(2012b) Guidance. CFA Society Toronto.

___________(2013) Monetary policy after the fall. Eric J. Hanson Memorial Lecture, University of Alberta.

Champagne, Julien, Guillaume Poulin-Bellisle, and Rodrigo Sekkel (2018) Evaluating the Bank of Canada staff economic projections using a new database of real-time data and forecasts. Bank of Canada Staff Working Paper 2018-52.

Faust, Jon and Jonathan H. Wright (2009) Comparing Greenbook and reduced form forecasts using a large realtime dataset. *Journal of Business and Economic Statistics* 27, 468–479.

Koeppl, Thorsten V. (2009) How flexible can inflation targeting be? Suggestions for the future of Canada's targeting regime. C.D. Howe Institute Commentary 293.

Kuttner, Kenneth N. (2004) The role of policy rules in inflation targeting. Federal Reserve Bank of St. Louis *Review* 86(4), 89–111.

Macklem, Tiff (2014) Flexible inflation targeting and 'good' and 'bad' disinflation. John Molson School of Business, Concordia University.

Poloz, Stephen S. (2013) Monetary policy as risk management. Canadian Club of Montréal.

Otto, Glenn D. and Graham M. Voss (2011) What do the RBA's forecasts imply about preferences over inflation and output volatility? *Economic Record* 87, 509–524.

___________(2014) Flexible inflation forecast targeting: Evidence from Canada. *Canadian Journal of Economics* 47, 398–421.

Rowe, Nicholas and James Yetman (2002) Identifying a policy-maker's target: An application to the Bank of Canada. *Canadian Journal of Economics* 35, 239–256.

Svensson, Lars E.O. (2003) What is wrong with Taylor rules? Using judgment in monetary policy through targeting rules. *Journal of Economic Literature* 41, 426–477.

____________(2009) Flexible inflation targeting — lessons from the financial crisis. Speech at the Netherlands Bank, 21 September 2009.

Woodford, Michael (2004) Inflation targeting and optimal monetary policy. Federal Reserve Bank of St. Louis *Review* 86, 15–41.

____________(2007a) The case for forecast targeting as a monetary policy strategy. *Journal of Economic Perspectives* 21(4), 3–24.

____________(2007b) Forecast targeting as a monetary policy strategy: Policy rules in practice. NBER working paper 13716.

Figure 1: *h*-Quarter-Ahead Inflation Forecasts

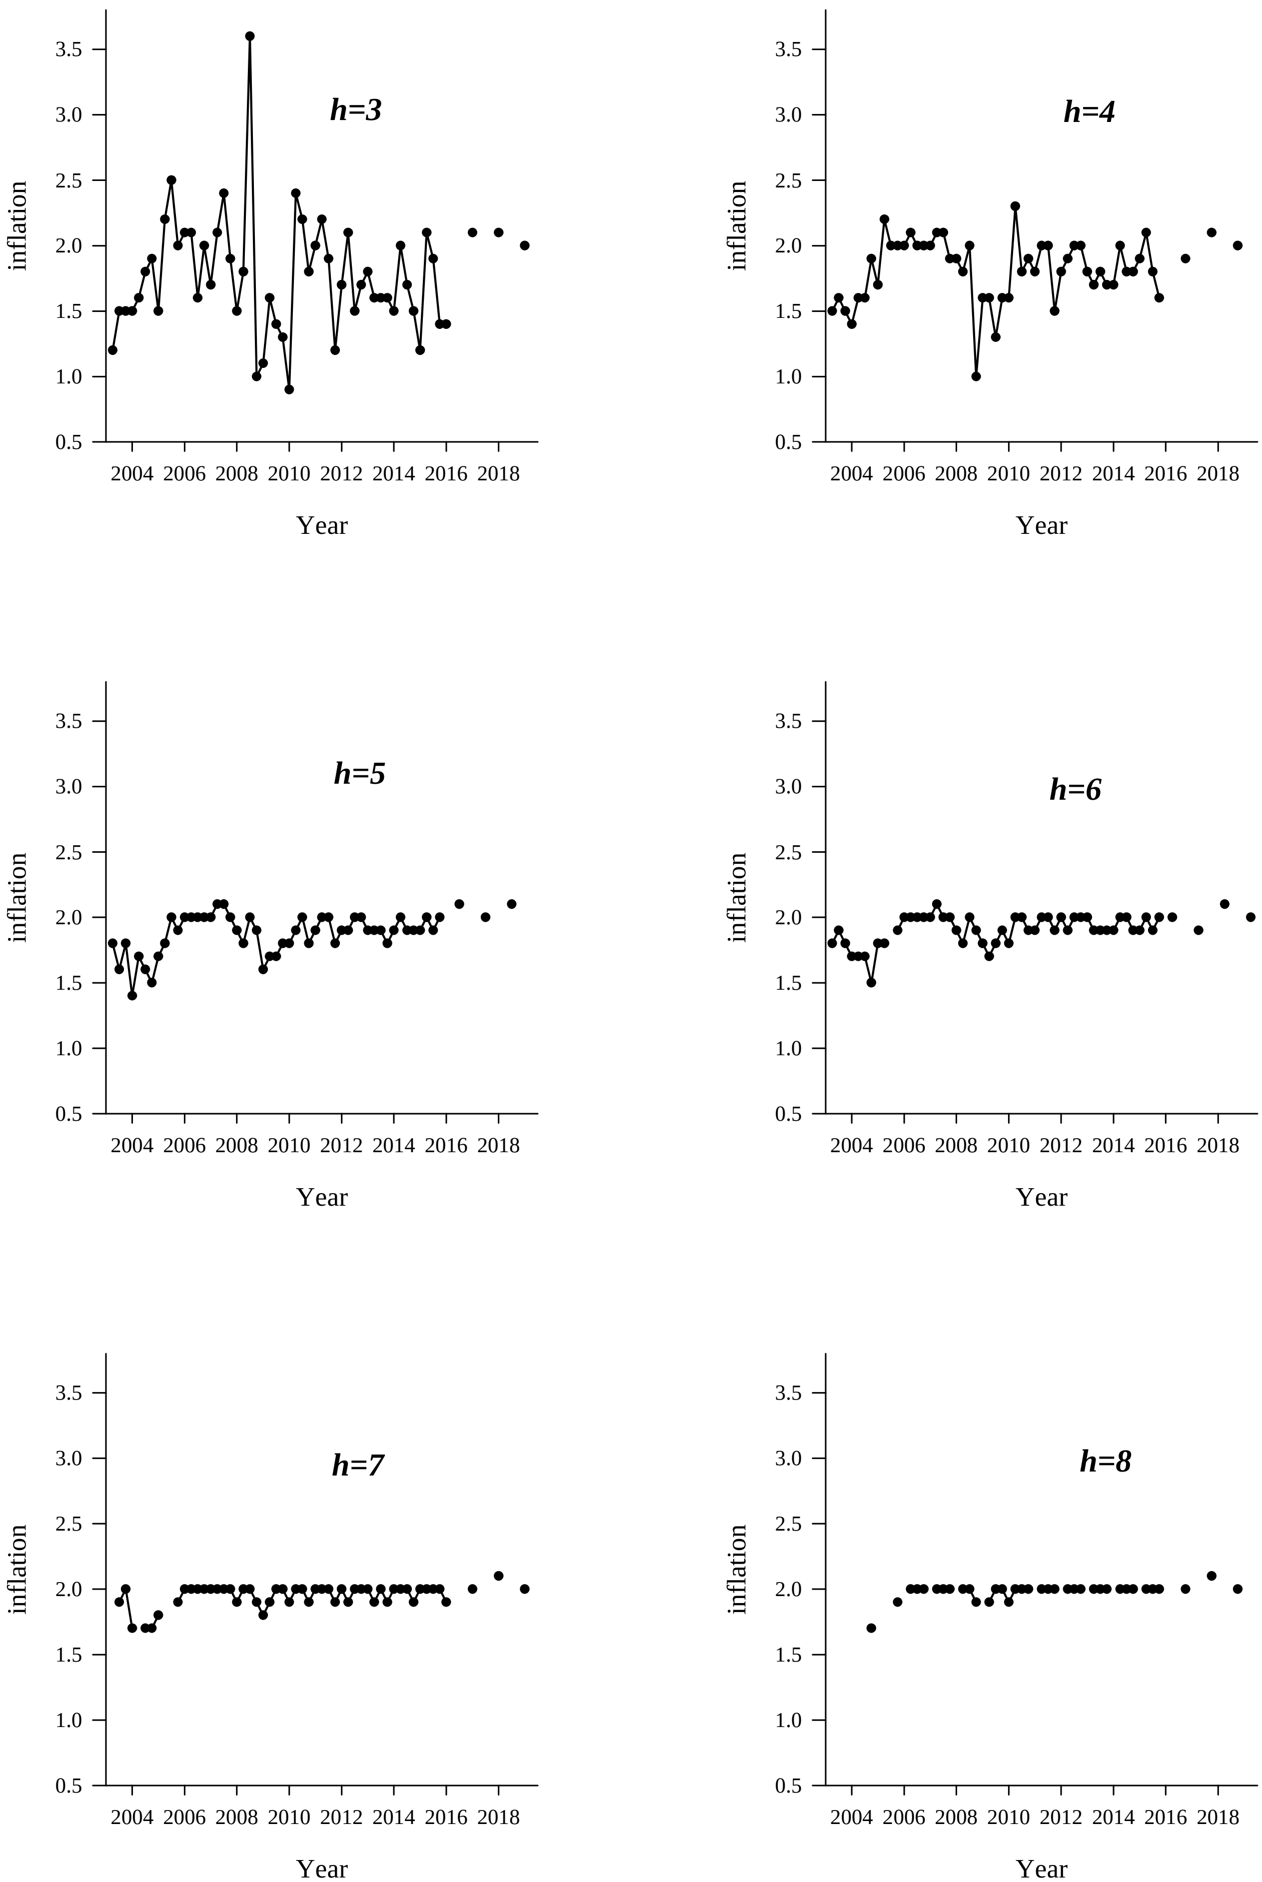

Table 1: Flexibility 1.0 Estimates and Tests
April 2003–April 2019

$$P_t \pi_{t+h} = \beta_0 + \beta_1 \pi_{t-1} + \epsilon_{t,h}$$

h	$\hat{\beta}_0$ (se)	$\hat{\beta}_1$ (se)	R^2	N
1	0.996 (0.181)	0.395 (0.084)	0.16	55
2	1.451 (0.095)	0.181 (0.054)	0.04	54
3	1.755 (0.136)	0.012 (0.071)	0.00	54
4	1.811 (0.091)	0.006 (0.48)	0.00	53
5	1.867 (0.048)	0.010 (0.020)	0.00	53
6	1.899 (0.033)	0.007 (0.012)	0.00	53
7	1.980 (0.012)	-0.014 (0.007)	0.02	50
8	2.004 (0.011)	-0.011 (0.006)	0.03	36

Table 1 (continued)

$$P_t\pi_{t+h} = \beta_0 + \beta_1\pi_t + \epsilon_{t,h}$$

h	$\hat{\beta}_0$ (se)	$\hat{\beta}_1$ (se)	R^2	N
1	0.325 (0.261)	0.783 (0.127)	0.61	56
2	0.773 (0.235)	0.556 (0.112)	0.33	55
3	1.395 (0.142)	0.210 (0.091)	0.16	55
4	1.717 (0.093)	0.057 (0.044)	0.04	54
5	1.825 (0.058)	0.031 (0.027)	0.03	54
6	1.880 (0.044)	0.017 (0.023)	0.02	54
7	1.966 (0.023)	-0.009 (0.015)	0.01	51
8	2.001 (0.026)	-0.010 (0.015)	0.02	36

Notes: h is the forecast horizon and N the number of observations. Standard errors are HAC Newey-West with $h-1$ lags.

Table 2: Flexibility 1.0 Estimates and Tests
July 2005–April 2019

$$P_t\pi_{t+h} = \beta_0 + \beta_1\pi_{t-1} + \beta_2 y_{t-1} + \epsilon_{t,h}$$

h	$\hat{\beta}_0$ (se)	$\hat{\beta}_1$ (se)	$\hat{\beta}_2$ (se)	R^2	N
1	0.952 (0.303)	0.198 (0.177)	0.219 (0.117)	0.26	47
2	1.413 (0.229)	-0.055 (0.175)	0.257 (0.093)	0.21	46
3	1.678 (0.091)	-0.058 (0.073)	0.132 (0.036)	0.16	46
4	1.773 (0.063)	-0.066 (0.049)	0.113 (0.020)	0.45	45
5	1.865 (0.022)	-0.005 (0.013)	0.041 (0.009)	0.32	45
6	1.906 (0.018)	-0.009 (0.012)	0.033 (0.006)	0.33	45
7	1.986 (0.016)	-0.024 (0.009)	0.016 (0.004)	0.17	45
8	2.011 (0.012)	-0.027 (0.009)	0.017 (0.003)	0.47	35

Table 2 (continued)

$$P_t\pi_{t+h} = \beta_0 + \beta_1\pi_t + \beta_2 y_t + \epsilon_{t,h}$$

h	$\hat{\beta}_0$ (se)	$\hat{\beta}_1$ (se)	$\hat{\beta}_2$ (se)	R^2	N
1	0.277 (0.318)	0.720 (0.150)	0.093 (0.098)	0.62	47
2	0.751 (0.305)	0.492 (0.145)	0.085 (0.103)	0.37	46
3	1.307 (0.161)	0.241 (0.132)	0.045 (0.053)	0.28	46
4	1.698 (0.047)	-0.023 (0.034)	0.109 (0.017)	0.46	45
5	1.823 (0.027)	0.022 (0.014)	0.037 (0.008)	0.39	45
6	1.866 (0.021)	0.020 (0.009)	0.025 (0.005)	0.37	45
7	1.968 (0.014)	-0.013 (0.008)	0.015 (0.004)	0.13	45
8	1.992 (0.010)	-0.013 (0.007)	0.013 (0.003)	0.21	35

Notes: h is the forecast horizon and N the number of observations. Standard errors are HAC Newey-West with $h-1$ lags.

Figure 2: *h*-Quarter-Ahead Output Growth Forecasts

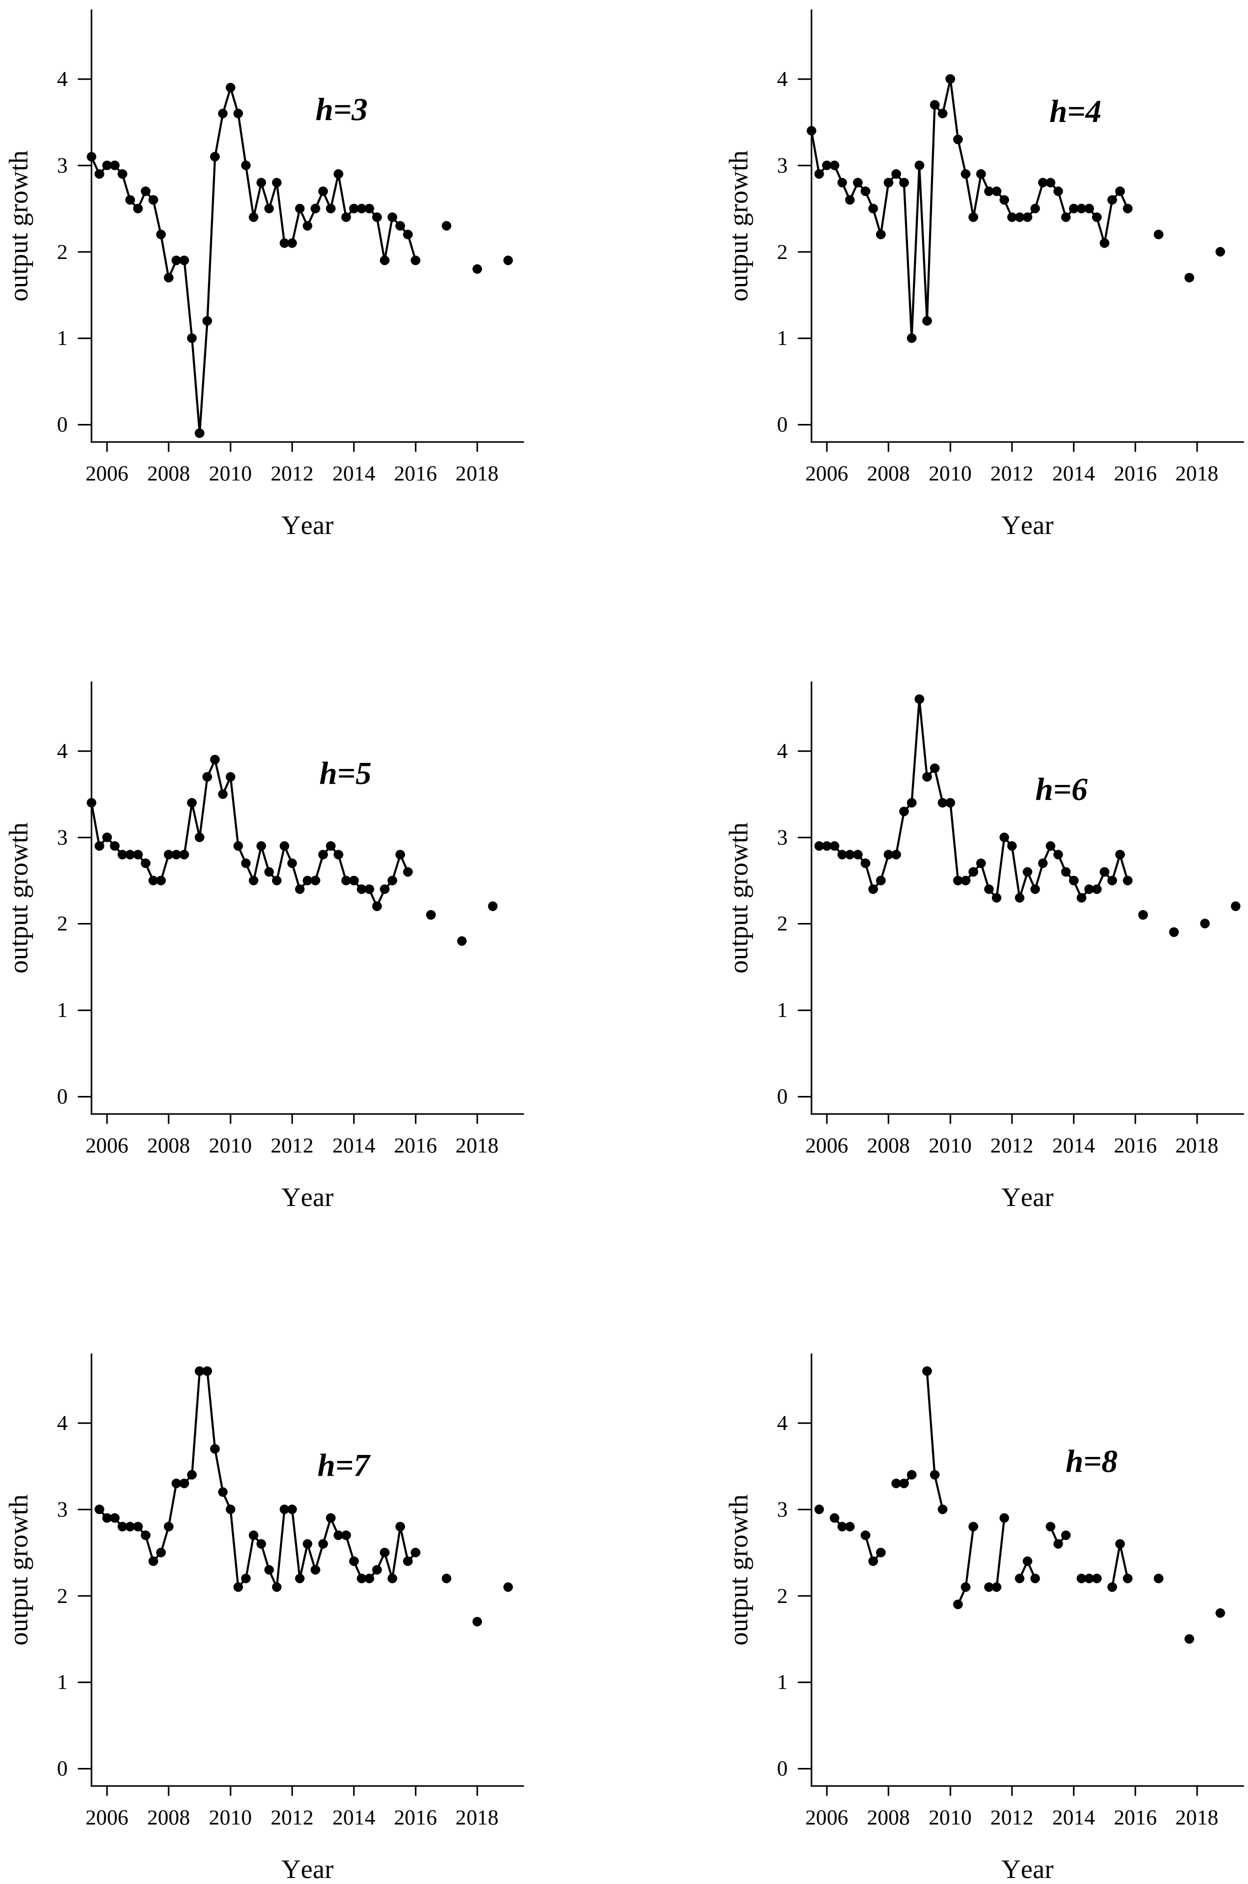

Table 3: Flexibility 2.0 Estimates and Tests
July 2005–April 2019

$$P_t\pi_{t+h} = \alpha - \phi P_t y_{t+h} + \epsilon_{t,h}$$

h	$\hat{\alpha}$ (se)	$\hat{\phi}$ (se)	$\hat{y}^*$ (se)	R^2	N
1	1.017 (0.146)	-0.342 (0.053)		0.29	47
2	1.201 (0.211)	-0.258 (0.113)		0.13	46
3	1.476 (0.230)	-0.134 (0.104)		0.04	46
4	1.706 (0.301)	-0.055 (0.111)		0.02	45
5	2.314 (0.097)	0.141 (0.036)	2.232 (0.164)	0.29	45
6	2.214 (0.051)	0.098 (0.019)	2.175 (0.145)	0.36	45
7	2.106 (0.027)	0.050 (0.009)	2.135 (0.162)	0.27	45
8	2.092 (0.022)	0.038 (0.008)	2.429 (0.099)	0.42	34
1–8	1.321 (0.146)	-0.217 (0.056)		0.15	353
4–8	2.053 (0.098)	0.044 (0.034)	1.203 (1.299)	0.03	214
5–8	2.169 (0.042)	0.078 (0.017)	2.154 (0.097)	0.26	169
6–8	2.138 (0.034)	0.063 (0.014)	2.178 (0.094)	0.29	124

Notes: h is the forecast horizon and N the number of observations. Standard errors are HAC Newey-West with 4 lags. The implicit target $\hat{y}^*$ is estimated as $(\hat{\alpha}-2)/\hat{\phi}$ for cases with $\hat{\phi}>0$.

Figure 3: Inflation and Output Growth

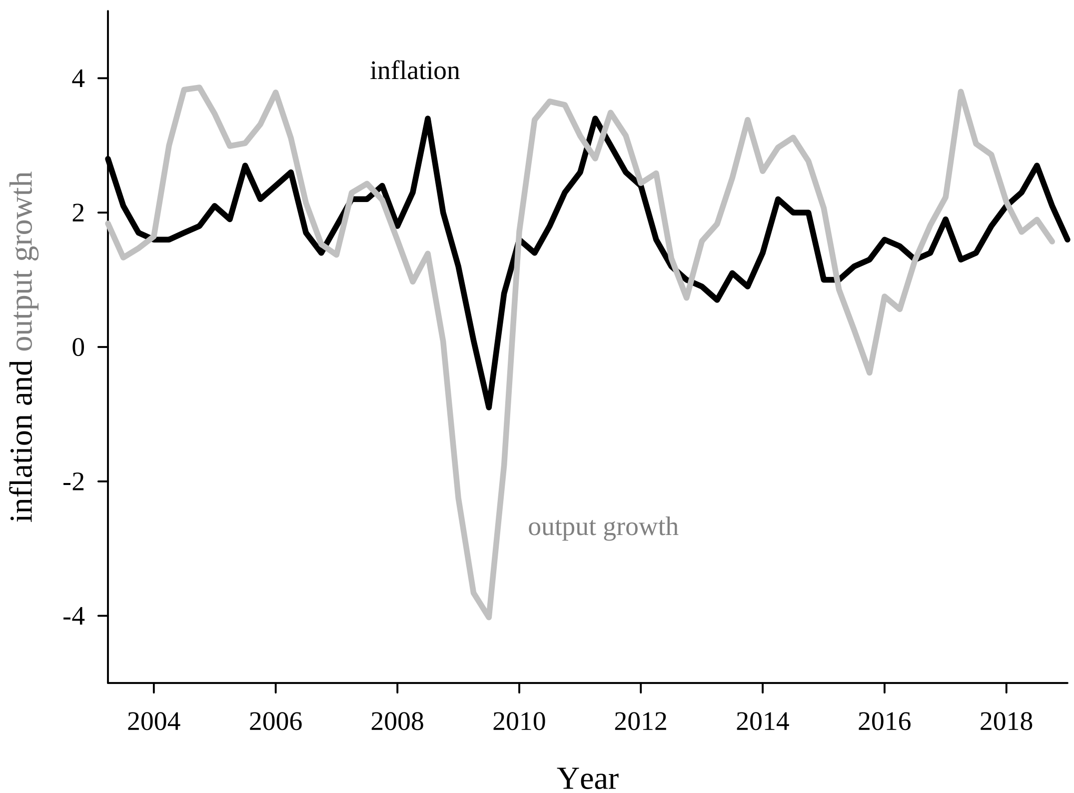

Note: Inflation is the 4-quarter growth rate in the CPI, from bankofcanada.ca. Output growth is the 4-quarter growth rate in real GDP, seasonally adjusted, from FRED.