University of Victoria | Economics

Department Discussion Paper DDP2012

ISSN 1914-2838

INFLATION TARGETS AND POLICY HORIZONS

Junhong Mu
Department of Economics, University of Victoria
Victoria, B.C., Canada V8W 2Y2
&
Graham M. Voss
Department of Economics, University of Victoria
Victoria, B.C., Canada V8W 2Y2

December 20 2022

Abstract

We use conditional forecasts from Canadian inflation over the inflation targeting period to estimate monetary policy horizons, defined as the period required for conditional forecasts of inflation to return to target. The conditional forecasts generate a set of policy horizons that tend to have median lengths longer than the typical two-year horizon identified by central banks. Further, the distribution of these forecasts is not unimodal, with policy horizons massed on shorter lengths of about one year and longer lengths of three or more years. The variation in the policy horizons, which are conditional on current economic conditions, is evidence of a form of flexibility in monetary policy. We also provide evidence of some instability in the underlying inflation models and associated policy horizons.

Keywords: monetary policy; inflation targeting; inflation

JEL Classifications: E31, E58

Author Contact:
Graham Voss, Dept. of Economics, University of Victoria, P.O. Box 1700, STN CSC, Victoria, B.C., Canada V8W 2Y2; E-mail: gvoss@uvic.ca

Inflation Targets and Policy Horizons

Junhong Mu
and
Graham M. Voss
University of Victoria, Canada

December 2022

Abstract

We use conditional forecasts from Canadian inflation over the inflation targeting period to estimate monetary policy horizons, defined as the period required for conditional forecasts of inflation to return to target. The conditional forecasts generate a set of policy horizons that tend to have median lengths longer than the typical two-year horizon identified by central banks. Further, the distribution of these forecasts is not unimodal, with policy horizons massed on shorter lengths of about one year and longer lengths of three or more years. The variation in the policy horizons, which are conditional on current economic conditions, is evidence of a form of flexibility in monetary policy. We also provide evidence of some instability in the underlying inflation models and associated policy horizons.

JEL Classification: E31, E58

Keywords: monetary policy; inflation targeting; inflation

Corresponding Author:

G Voss
Department of Economics
University of Victoria

Email: gvoss@uvic.ca

This paper has benefited from discussions with Vasco Gabriel, Michael King, Glenn Otto, and Felix Pretis. The authors are responsible for any errors and omissions.

1 Introduction

Inflation targeting central banks operate with a policy horizon, the period over which monetary policy actions are focused and by the end of which, or before, inflation is expected to meet the inflation target. Framing monetary policy decisions and expected outcomes within a policy horizon is an important part of a transparent inflation targeting framework and central bank communication. For most inflation targeting central banks, the horizon is generally somewhere around two years, though often with various qualifications. For Canada, which is the focus of this study, the Bank of Canada currently identifies a horizon of typically six to eight quarters to return inflation to target.[1] Statements by central banks about policy horizons may not, however, reflect the outcomes we observe. This paper proposes a simple method of estimating the *de facto* horizons for monetary policy and applies these methods to Canada over the inflation targeting period, which began in the early 1990s.

Our measures of policy horizons are based on conditional forecasts from statistical models for inflation. The underlying logic is that the monetary policy actions of the Bank of Canada within the inflation targeting framework, in conjunction with the rest of the economy, determine the stochastic process for inflation. In the next section, we provide a very simple model to motivate this reasoning. With a suitably estimated model for inflation, we can then use conditional forecasts to identify the time periods required for the forecasts to return to target. We can then examine the distribution of the time periods — which are date-specific policy horizons — to determine the typical policy horizon and the range of policy horizons observed.

We use these methods for two measures of inflation for Canada: the All-items CPI and a measure of core inflation. The first is the measure of inflation targeted by the Bank of Canada while the second is a measure of underlying inflation that may better reflect monetary policy objectives. In the first instance, we focus on models estimated over the full inflation targeting period. We then consider models estimated over different sub-samples, to see whether the results are stable over time. For a given sample, our methods provide a distribution of policy horizons, one for each date of the sample, which characterizes monetary policy. To foreshadow our main results, the distributions tend to have median policy horizons longer than the typical two-year horizon identified by central banks. More importantly, the distributions are not unimodal, with horizons massed on shorter lengths of about one year and longer lengths of about three years. This is evidence of a certain type of flexibility, which we explain below, and suggests that the concept of a typical policy horizon, as used in central bank communications, obscures the variation of measured policy horizons.

The paper proceeds as follows. In section 2, we discuss more generally the issues of monetary policy horizons and the related literature. In section 3, we present our empirical framework, including a simple model of monetary policy that allows us to fix ideas about how we measure policy horizons. Section 4 presents our empirical results and section 5 concludes and provides suggestions for further work.

2 Policy Horizons — A Brief Review

Inflation targeting by central banks is best interpreted as inflation forecast targeting. As emphasized in Svensson (1999), the central banks's conditional inflation forecast — and possibly those of other target variables — serve as the intermediate target for monetary policy. In practice, an inflation forecast target

[1] Bank of Canada (2022).

requires a timeframe, or horizon, over which the conditional forecast is brought toward target. King (1994, p. 118) provides a nice, early interpretation of inflation targeting along these lines:[2]

> The use of an inflation target does not mean that there is no intermediate target. Rather, the intermediate target is the expected level of inflation at some future date chosen to allow for the lag between changes in interest rates and the resulting changes in inflation. In practice, we use a forecasting horizon of two years.

Once inflation targeting is interpreted as targeting forecasted inflation (or other target variables), the need for a policy horizon is immediate. Nonetheless, it is useful to clarify what determines the policy horizon and the role it has in monetary policy.

First, monetary policy operations affect target variables with lags; indeed, this is the principal reason why inflation targeting is best interpreted as using conditional forecasts as intermediate targets. The policy horizon from this perspective identifies a period by which monetary policy can reasonably be expected to return inflsation to target given current conditions.[3] As Batini and Nelson (2001a) discuss, a consensus emerged over the 1990s, when a number of countries adopted inflation targeting, of a two-year lag between monetary policy actions and their main effect on inflation.[4]

The second reason for specifying a policy horizon is to provide transparency about the objectives of the central bank in terms of how quickly it seeks to return inflation to target. From this perspective, the emphasis is not when monetary policy can be effective but rather how long monetary policy will allow for the adjustment of forecasted inflation to target. This second role for policy is important for communicating the commitment of the central bank to the inflation target and serves to anchor inflation expectations. From this perspective (and conditioning on whatever lags of monetary policy exist), the choice of policy horizon reflects a trade-off, as emphasized in Coletti et al. (2006) in its discussion of the Bank of Canada's policy horizon. Too short a horizon risks destabilizing the economy through excessive volatility in the financial and real economy caused by aggressive changes in monetary policy, while too long a horizon risks persistent deviations of inflation from target (and, by implication, a loss of credibility).

More recently, policy horizons are used by many inflation targeting central banks as a means of introducing, and communicating, flexibility in their monetary policy actions and objectives.[5] For example, the Bank of Canada when describing its framework states (emphasis in original):[6]

> Canada's inflation-targeting framework is also flexible. Typically, the Bank seeks to return inflation to target over a horizon of six to eight quarters. However, the most appropriate horizon

[2]At the time, Mervyn King was the chief economist at the Bank of England. Svensson (1999) quotes this as an early statement by a practicing central banker to highlight the role of conditional forecasts. Of interest here as well is the clear statement about a policy horizon and the identification of the two-year horizon commonly used by central banks.

[3]The lags of monetary policy were first emphasized and measured in Friedman (1961) and Friedman (1972). More recently, Batini and Nelson (2001a) reconsiders the empirical evidence and finds similar results.

[4]The authors argue that this is consistent with the estimates of Friedman (1972) as well as their own, updated estimates. Their conclusion on a consensus of two-years is based primarily on the discussion in Bernanke et al. (1999), though they discuss other studies that provide support, directly or indirectly, for horizons of this length. The two-year horizon is also consistent with King (1994) discussed above.

[5]See Mu (2022) for a survey of OECD central banks and their use of policy horizons as a means of implementing flexibility in inflation targeting. See also Smith (2019), which we discuss further below, for a recent examination of flexibility in the Bank of Canada's monetary policy.

[6]Bank of Canada (2022, preamble). The preamble entitled 'Canada's inflation-control strategy' is the most comprehensive statement by the Bank on how it implements the Monetary Policy Framework Agreement with the Government of Canada.

for returning inflation to target will vary depending on the nature and persistence of the shocks buffeting the economy.

For the Bank of England, the use of policy horizons in this manner is an explicit requirement under its remit from the UK Treasury (emphasis added):[7]

> In forming and communicating its judgements the Committee should promote understanding of the trade-offs inherent in setting monetary policy to meet a forward-looking inflation target while giving due consideration to output volatility. It should set out in its communication: ... *the horizon over which the Committee judges it is appropriate to return inflation to the target; the trade-off that has been made with regard to inflation and output variability in determining the scale and duration of any expected deviation of inflation from the target; ...*

Along with the practical policy communications for policy horizons are a number of theoretical and empirical studies that examine policy horizons within the inflation-targeting context. Most notably, Batini and Nelson (2001b) is an early paper that provides both theoretical context as well as empirical assessment of policy horizons.

Batini and Nelson (2001b) distinguish between two types of policy horizons. The first is termed optimal policy horizons and corresponds to the policy horizons we examine this paper. The second is termed optimal feedback horizons and relates to forward-looking policy rules. In this second case, the concern is the appropriate horizon for projections of target variables as used in policy rules. As Batini and Nelson (2001b) note, these are different concepts and not immediately related.[8] The objective of Batini and Nelson (2001b) is to calculate optimal horizons (of both types) for inflation targeting using estimated or calibrated structural models for the UK economy. For the optimal policy horizon, of interest here, the paper specifies a loss function for the central bank and then uses the models to determine the optimal paths for the economy. Implicit in these solutions are response functions to different shocks from which optimal policy horizons can be determined for each shock. In this way, the authors determine the period of time for which it is least costly to bring inflation back to target after a shock. The optimal policy horizons they identify vary depending upon the shock and the underlying model, ranging from 8–19 quarters.

Two similar studies have been done for Canada, Cayen et al. (2006) and Basant-Roi and Mendes (2007). These are nicely summarized in Coletti et al. (2006). Both Cayen et al. (2006) and Basant-Roi and Mendes (2007) use structural models of the Canadian economy with a specified loss function over inflation, the output gap, and variations in interest rates — similar to Batini and Nelson (2001b).[9] Both these papers focus solely on optimal policy horizons (rather than optimal feedback horizons) and both find similar results. An important distinction between these two studies and Batini and Nelson (2001b) is that the horizons are calculated not for a single shock but from stochastic simulations of all shocks in the respective models over a given sample period. This gives rise to a distribution of policy horizons and is conceptually very similar to our own empirical methods. The mean policy horizons from the two studies are similar at six or seven quarters, and the distribution of horizons, which range considerably, are similar as well. The mean policy

[7] Hunt (2022); available at www.bankofengland.co.uk.

[8] For example, Giannoni and Woodford (2003) demonstrate that optimal policy rules need not put much or any weight on target variable forecasts even when inflation inertia, which introduces lags to monetary policy, is assumed. See also Batini et al. (2006) and Gabriel et al. (2009) for studies that examine policy horizons from the perspective of forward-looking policy rules.

[9] Cayen et al. (2006) uses a large structural model that is used regularly by the Bank of Canada for policy analysis. Basant-Roi and Mendes (2007) uses a smaller model and examines on the implications of housing-price bubbles for policy horizons.

horizons are consistent with the Bank's statement on policy horizons discussed above, with the mean value the basis for the Bank's use of the term 'typically' when describing its policy horizons.[10]

A feature of these studies (Batini and Nelson (2001b), Cayen et al. (2006) and Basant-Roi and Mendes (2007)), all of which define policy horizons in terms of the convergence of inflation to its target, is that for the dynamic models used, the target variable will only approach the target asymptotically. Consequently, it is necessary to define the policy horizon in terms of when a target variable is within a certain tolerance of the target.[11] This is discussed further in the empirical section and has a considerable bearing on our results.

A somewhat related technical issue is whether to measure the policy horizon in absolute or relative terms. The absolute measure calculates the number of periods for the target variable to return to target (within the tolerance band). The relative measure calculates the number of periods for the target variable to eliminate a certain percentage of the original deviation. Here, we follow Cayen et al. (2006) and Basant-Roi and Mendes (2007) and focus on the absolute measure. In our view, central banks are more likely to communicate and focus on policy horizons that are defined in terms of returning to the target, rather than a set percentage of an historical deviation.

There are a number of important differences between our approach and these studies. First, ours is purely empirical with no reference to any underlying structural model. In contrast, these studies identify optimal policy horizons subject to some loss function. In other words, our objective is to measure the policy horizons implicit in central bank behaviour, as captured within the process for inflation, rather than estimate or model the optimal policy horizon.

Second, and relatedly, we focus on policy horizons for conditional forecasts rather than those for impulse response functions of the target variables themselves. This latter distinction is quite important. It means, firstly, that we are not identifying shock-specific policy horizons as done in Batini and Nelson (2001b). In this regard, we are closer to Cayen et al. (2006) and Basant-Roi and Mendes (2007), which use stochastic simulations across a broad set of shocks. However, whether a single shock is used or multiple shocks through stochastic simulations, the policy horizons are not state dependent in that they do not depend upon current and recent past levels of inflation or other variables of interest. In contrast, our policy horizons are state dependent; the length of time it takes to return to target for a conditional forecast for inflation will depend upon current and recent past values of inflation.[12]

A further important distinction is that these studies of optimal policy horizons are able to identify optimal policy horizons that are implicitly flexible — horizons that differ depending upon the nature of the shocks, given the loss function for the central bank and the modelled structure of the economy. In contrast, our conditional-forecast policy horizons arise from the estimated path for inflation and vary only because of the state of the economy as captured in the variables underlying the inflation forecast. We discuss this in further detail once we have provided more structure to our approach.

In addition to these studies of policy horizons, there is a related empirical literature based on the inflation forecast targeting approach of Svensson (1999) that also provides information on flexibility and policy horizons. Otto and Voss (2014), building on earlier work by Rowe and Yetman (2002), estimates flexible

[10] The policy horizons from Basant-Roi and Mendes (2007) when housing-price bubbles are included are considerably longer and have a wider distribution than those without. The results from Cayen et al. (2006) are based on the Bank's primary policy model and are arguably most relevant.

[11] This is also relevant for a more recent study, Davig and Foerster (2021), discussed further below.

[12] Our empirical analysis focuses exclusively on inflation. The methods extend, in a straightforward manner, to a multivariate analysis, in which case the policy horizons would depend upon a broader set of conditioning information.

inflation targets for Canada, targets that explicitly weight both inflation and some measure of economic activity. This relates to earlier notions of flexible inflation targets as discussed in Svensson (1999), rather than flexibility enacted through changes in the policy horizon.[13] The methods used, however, do provide measures of average policy horizons conditional on the estimated flexible targets. For example, Otto and Voss (2014) finds evidence of a flexible target for the Bank of Canada and an average policy horizon of 12–18 months, consistent with the Bank of Canada's typical horizon.

A second strand of this literature examines flexibility in inflation targets by examining the forecasts of the central banks themselves, rather than estimating conditional forecasts from the data. Kuttner (2004) is the first study to pursue this, followed by Otto and Voss (2011) for Australia, and more recently Smith (2019) for Canada. The latter is the most closely related to our objectives and methods as it is designed to see if flexibility in monetary policy has bearing on the policy horizons of the Bank of Canada.

Smith (2019) proceeds by collecting published projections for inflation and output growth from the Bank's *Monetary Policy Reports*. The study then examines whether the projections for inflation deviate from the inflation target of two percent at horizons of up to eight quarters, and if these deviations can be explained by current conditions, either inflation alone or inflation and output growth (in both cases, current or single-lagged values are used). The results suggest that even at longer horizons (7–8 quarters), current and lagged inflation and output growth do predict deviations from target for projected inflation, suggesting that the Bank does exercise some flexibility through variations in the policy horizon; that is, depending upon current conditions, the Bank does not always project a return to target by the end of eight quarters. However, as the study emphasizes, while the contributions are statistically significant they are economically quite small.[14] In other words, what flexibility there may appear to be is not substantially altering the policy horizon for the Bank.

Our approach has the same motivation as Smith (2019) but differs in a number of substantial ways. First, we use conditional forecasts taken from the data, which should summarize the actual behaviour of the Bank of Canada (and the underlying economic environment) rather than their reported projections. Second, we focus on measuring the policy horizon directly, rather than indirectly based on the distance from target at horizons over the typical two-year projection horizon in Bank publications. A limitation of our study, relative to Smith (2019), is that we summarize current conditions — the basis for the inflation forecasts — purely in terms of the recent history of inflation, rather than other influences such as output growth. Our methods, though, readily extend in this direction. Despite the differences in approach, the two studies complement each other. The focus on central bank projections provides direct information about how the Bank uses flexibility, particularly how it communicates that flexibility through its projections. The focus on conditional forecasts based on the actual inflation process provides evidence on what the Bank has been able to achieve and the variation in policy horizons due to current conditions.

Finally, Davig and Foerster (2021) is a recent study that further highlights the important practical role of policy horizons for inflation targeting. These authors demonstrate, within standard theoretical frameworks, that a central bank can achieve an optimal rule-based policy through explicit communication of target horizons for target variables. The principle is relatively straightforward. Within a standard modelling framework, a policy rule implies a specific path for future inflation (and other target variables if relevant)

[13] See also Woodford (2007) for a discussion of flexible inflation targets in practice.

[14] Smith (2019) also considers a flexible inflation target, similar to Otto and Voss (2011); that is, a target with both inflation and output growth. Again, while there is evidence that deviations from the compound target at longer horizons are correlated with current conditions, suggesting the policy horizons are flexible, the magnitude of the effect is very small.

converging to target at some point based upon the shocks to the economy. By the central bank credibly communicating the path (in effect, the policy horizon), it is in effect committing to an underlying policy rule. From our perspective, this study nicely captures the link between monetary policy, the economy, and policy horizons. In the following section, we present a simple theoretical example, based on that in Davig and Foerster (2021), to motivate and understand the empirical analysis that follows.

3 Empirical Framework

3.1 A Theoretical Example

The empirical methods used in this paper are non-structural; nonetheless, it is useful to motivate our approach with a standard model of monetary policy. For this, we use the three equation New Keynesian (NK) model as presented in Davig and Foerster (2021). The authors use this model to motivate the relationship between effective communication of (conditional) policy horizons combined with inflation target tolerance bands and strict monetary policy rules. Many of the issues these authors address are also relevant for this study and the NK model is a nice means to motivate our empirical work and to identify a number of relevant issues.

The model consists of the usual three behavioural equations: a New Keynesian IS curve (the consumption Euler equation in equilibrium); an aggregate supply condition; and a simple rule to describe monetary policy. Each of these equations is log-linearized around the steady-state. The equations are:[15]

$$\tilde{x}_t = E_t\tilde{x}_{t+1} - \sigma^{-1}(\tilde{i}_t - E_t\pi_{t+1}) + g_t \tag{1}$$

$$\tilde{\pi}_t = \beta E_t\tilde{\pi}_{t+1} + \kappa\tilde{x}_t + u_t \tag{2}$$

$$\tilde{i}_t = \alpha\tilde{\pi}_t + \eta_t \tag{3}$$

The endogenous variables are: $\tilde{x}_t$, the period t output gap (deviations of output from the natural level of output); $\tilde{i}_t$, the nominal interest rate, expressed as deviations from the steady-state nominal interest rate; and $\tilde{\pi}_t$, the inflation rate, expressed as deviations from steady-state inflation, which is the inflation target π^*. The exogenous shocks are g_t, an aggregate demand shock; u_t an aggregate supply shock; and η_t, a policy shock. The parameters are standard and we assume $\alpha > 1$ to satisfy the Taylor principle.

The aggregate demand and supply shocks follow:

$$g_t = \rho_g g_{t-1} + \epsilon_{g,t} \tag{4}$$

$$u_t = \rho_u u_{t-1} + \epsilon_{u,t} \tag{5}$$

with $0 \leq \rho_g < 1$ and $0 \leq \rho_u < 1$. The three innovations, $\{\epsilon_{g,t}, \epsilon_{u,t}, \eta_t\}$ are i.i.d. with mean zero and variances $\{\sigma^2_{\epsilon_g}, \sigma^2_{\epsilon_u}, \sigma^2_{\eta}\}$. The three innovations are assumed to be independent of each other. Finally, given

[15] The notation here differs slightly from Davig and Foerster (2021). Here we emphasize the deviations from steady-state, which allows us to map the discussion more closely to the empirical work.

the processes for the two shocks, we have

$$\sigma_g^2 = \frac{\sigma_{\epsilon_g}^2}{1-\rho_g} \qquad \sigma_u^2 = \frac{\sigma_{\epsilon_u}^2}{1-\rho_u}$$

The solution (from Davig and Foerster (2021)) for inflation, expressed as deviations from steady-state, is:[16]

$$\tilde{\pi}_t = \Psi_1 g_t + \Psi_2 u_t + \Psi_3 \eta_t \tag{6}$$

Or, for inflation itself:

$$\pi_t = \pi^* + \Psi_1 g_t + \Psi_2 u_t + \Psi_3 \eta_t \tag{7}$$

where

$$\Psi_1 = \frac{\kappa}{1+\alpha\kappa\sigma^{-1}-\Phi_g} \quad \Psi_2 = \frac{1-\rho_u}{1+\alpha\kappa\sigma^{-1}-\Phi_u} \quad \Psi_3 = \frac{\kappa\sigma^{-1}}{1-\alpha\kappa\sigma^{-1}}$$

and,

$$\Phi_z = \rho_z(1+\kappa\sigma^{-1}+\beta(1-\rho_z)) \qquad z = g, u \tag{8}$$

Our interest is in the stochastic properties of inflation, π_t. From equation (7), inflation has a mean of π^* and a variance of

$$\begin{aligned} \sigma_\pi^2 &= \Psi_1^2\sigma_g^2 + \Psi_2^2\sigma_u^2 + \Psi_3^2\sigma_\eta^2 \\ &= \Psi_1^2\Big(\frac{\sigma_{\epsilon_g}^2}{1-\rho_g}\Big) + \Psi_2^2\Big(\frac{\sigma_{\epsilon_u}^2}{1-\rho_u}\Big) + \Psi_3^2\sigma_\eta^2 \end{aligned} \tag{9}$$

For purposes of discussion, it is useful to group parameters into policy parameters — $\Pi_p = \{\alpha, \pi^*, \sigma_\eta^2\}$; and structural economic parameters — $\Pi_s = \{\kappa, \sigma, \rho_g, \rho_u, \sigma_{\epsilon_g}^2, \sigma_{\epsilon_u}^2\}$. In doing so, we can summarize the variance of inflation as $\sigma_\pi^2 = \sigma_\pi^2(\Pi_p, \Pi_s)$, which emphasizes its dependence upon both the underlying policy parameters as well as the structure of the underlying economy.

The dynamic process for inflation is:

$$\begin{aligned} \pi_t &= \pi^* + \Psi_1 g_t + \Psi_2 u_t + \Psi_3 \eta_t \\ &= \pi^* + \Psi_1(1-\rho_g L)^{-1}\epsilon_{g,t} + \Psi_2(1-\rho_u L)^{-1}\epsilon_{u,t} + \Psi_3\eta_t \end{aligned}$$

Collecting terms,

$$\begin{aligned} (1-\rho_g L)(1-\rho_u L)\pi_t = (1-\rho_g L)(1-\rho_u L)\pi^* + \Psi_1(1-\rho_u L)\epsilon_{g,t} + \\ \Psi_2(1-\rho_g L)\epsilon_{u,t} + \Psi_3(1-\rho_g L)(1-\rho_u L)\eta_t \end{aligned} \tag{10}$$

This is clearly a complex ARMA process for inflation, involving the two underlying structural shocks and

[16] Solutions for the other endogenous variables are not presented here as they are not required for the discussion here. See Davig and Foerster (2021) for the full solution.

the policy innovation. Further, as anticipated, the inflation process is a function of the model parameters, $\{\Pi_p, \Pi_s\}$.

The final part of the inflation process of interest is the conditional forecast, which is immediately evident from equation (7):

$$E_t\pi_{t+j} = \pi^* + \Psi_1\rho_g^j g_t + \Psi_2\rho_u^j u_t \tag{11}$$

As emphasized by Davig and Foerster (2021), these conditional forecasts — which are the basis for the policy horizons for inflation forecast targeting — are functions of both the structural and the policy parameters. In their analysis, they examine policy horizons using these conditional forecasts for different shocks; these are equivalent to the impulse response functions for the model.

For our empirical analysis, however, we are not able to observe the underlying shocks so our approach is somewhat different. Even if we were to use a structural model to identify the underlying shocks, it is not obvious that this is really how central banks use and communicate policy horizons. Rather, our focus is on the state of the economy as measured by observable variables. In the current context, we limit this focus to the current and past levels of inflation.

The first step is to assume that for our structural model of inflation above, the process for inflation can be well-approximated by an ARMA(p,q) process.[17] Specifically, we assume that equation (10) can be well-approximated as follows:

$$\rho(L)(\pi_t - \pi^*) = \Theta(L)\epsilon_t \tag{12}$$

where $\rho(L) = 1-\rho_1 L-\rho_2 L^2-\ldots-\rho_p L^p$ and $\Theta(L) = (1-\theta_1 L-\theta_2 L^2-\ldots-\theta_q L^q)$. Under these assumptions, the conditional expectations can be written as linear functions of current and lagged inflation and MA error terms:

$$\begin{aligned} E_t(\pi_{t+j} - \pi^*) &= f(\pi_t, \ldots, \pi_{t-p}, \epsilon_t, \ldots, \epsilon_{t-q}; \rho, \Theta) \\ &= f(\chi_t; \rho, \Theta) \end{aligned} \tag{13}$$

where ρ and Θ represent the set of parameters of the AR and MA components respectively, and χ_t denotes the conditioning information set for the projections. The reduced form in equation (12) can be estimated and used to construct the conditional forecasts in equation (13).

As emphasized by Batini and Nelson (2001b) and Davig and Foerster (2021), the horizon for inflation to return to zero (the target) is infinite. To make the policy horizon meaningful requires specifying a tolerance band, which we denote μ. Next, we denote a conditional policy horizon, N_t, which is the time taken to return expected inflation to within μ of the target. Since there is no reason to believe that the conditional forecasts will approach the target monotonically, we need to define the target in terms of absolute deviations. The policy horizon N_t is then implicitly defined as follows:

$$|E_t\pi_{t+j} - \pi^*| < \mu \qquad \forall j \geq N_t(\chi_t; \rho, \Theta; \mu) \tag{14}$$

[17]By assumption, our model is stationary so there is no need for an ARIMA process. As well, we could assume an AR(p) process; however, as our empirical work finds an ARMA process necessary to model inflation, it seems reasonable to work with that framework here as well.

Here we have noted the dependence of N_t on the choice of tolerance level μ, as well as the underlying structural and policy parameters and the conditioning information set.[18]

We are now in a position to discuss various features of the policy horizons defined in equation (14). These features directly bear on how we interpret our estimated policy horizons below.

First, the policy horizons depend critically upon the choice of μ. Intuitively, a smaller μ will give rise to longer policy horizons. Our empirical analysis will explore a range of choices for μ but as a concrete example, Figure 1 identifies the policy horizons for two tolerance levels, $\mu = 0.10$ and $\mu = 0.25$, for a simple AR(2) model. Because this model has non-monotonic forecasts, it demonstrates the need for policy horizons to be defined in terms of absolute value deviations.[19] From the figure, we see that for $\mu = 0.10$, $N = 16$ and for $\mu = 0.25$, $N = 10$.

Second, the policy horizons are a function of time, through the conditioning variables for the forecasts, χ_t. So, for a given sample and realizations of χ_t, there will be a set of policy horizons with their own distribution. Our empirical analysis examines the distributions for our data sample. In this respect, our analysis is very similar to Cayen et al. (2006) and Basant-Roi and Mendes (2007), both of which examine the distribution of policy horizons arising from stochastic simulations of their models.[20]

Third, the policy horizons here are *flexible* but in a very specific and narrow manner. The policy horizons will vary systematically with the state of the economy, χ_t. We conjecture that, in general, the further the economy is away from its steady state (e.g., current and recent history of inflation above target), the longer will be the policy horizon. While we are unable to prove this in the general environment here, this seems intuitive; moreover, it is very consistent with statements from the Bank of Canada.[21]

While this is a narrow type of flexibility, it is nonetheless an important aspect of the monetary policy environment. One could imagine, for example, a central bank that sets policy to ensure that inflation is back on target irrespective of the current level of inflation and implications for other parts of the economy.[22] A further point to emphasize is that the behaviour of the central bank and the structure of the economy over time determines the process for inflation, so there is no sense in which one can link an estimated policy horizon at a point in time with other variables outside of χ_t at that point in time.[23]

Fourth, as these are forecasts, it is very natural to interpret them from this perspective. This raises questions about whether they are good forecasts. A further concern that always arises when discussing forecasts is in-sample versus out-of-sample. There is, however, an alternative perspective that better describes our approach and objectives. Given a stable model for inflation — the combination of the central bank's behaviour and the economy — the policy horizon estimates can be thought of as summary statistic of monetary policy, one that is conditional on the current state χ_t. In this respect, it is comparable to historical error decompositions of the endogenous variables of a structural VAR model.

[18] As noted earlier, we are setting the tolerance in absolute terms rather than relative terms.

[19] For completeness, the model is: $\pi_{t+1} = (1 - \rho_1 - \rho_2)\pi^* + \rho_1\pi_t + \rho_2\pi_{t-1} + \epsilon_{t+1}; \rho_1 = 0.7; \rho_2 = -0.7$. The projections are conditional on $\pi_t = 3.0$ and $\pi_{t-1} = 1.00$.

[20] In contrast, Batini and Nelson (2001b) and Davig and Foerster (2021) focus on policy horizons specific to a single realized shock; in effect, the impulse response functions of their structural models. The different emphasis here reflects the different nature of our study.

[21] Davig and Foerster (2021) demonstrate that this is the case in their analysis, for a calibrated version of the model above, by varying the size of the shock being examined.

[22] See Mu (2022) for a discussion of how different central banks implement policy horizons, including those with strict policy horizons.

[23] We return to this issue after we have presented the empirical analysis, highlighting the scope of our analysis and suggestions for future directions.

Fifth, there are two ways in which the policy parameter α affects the policy horizon. First, there is an effect directly through the dynamics of the inflation process, as evident in equation (7) and the dynamic specifications that follow. For example, a larger α will dampen the effects of the aggregate demand and supply shocks, thus shortening the policy horizon. The second comes from the effect of α on the distribution of inflation. From equation (9), a larger α will tend to reduce the variance of inflation; in other words, monetary policy will ensure that inflation is closer to target on average. This has implications for the distribution of the policy horizons. For simplicity, assume that χ_t is only current inflation, π_t. A larger α that generates a distribution of π_t that is closer to target (smaller variance) will give rise to shorter policy horizons; that is, it will reduce the mean (median) policy horizon over the history of monetary policy.

A similar conclusion applies to reducing the variance of the monetary policy innovations and in this case, we can loosely interpret this as a linkage to the flexibility of monetary policy. One way to interpret flexible monetary policy is that occasionally, the central bank may depart from its usual behaviour because of current conditions in the economy. Greater flexibility in this regard can be thought of as a higher variance for η_t, the policy innovation. This gives rise to higher variance for inflation and by implication longer mean (median) policy horizons.

Finally, the model structure and the definition of the flexible policy horizon clearly indicate that the forces underlying the policy horizon estimates and distribution are functions of many unobservable parameters, which are either structural and policy-related. The analysis above assumes that these are stable. In the empirical analysis, we investigate the stability of our policy horizons. The solutions above also indicate that any evidence of instability will be very difficult to allocate to one of either changes in policy or changes in the underlying economic structure.

4 Estimation

4.1 Data

The focus of this paper is the behaviour of inflation — and the associated policy horizons — under the Bank of Canada's inflation targeting framework. This framework was announced jointly by the Bank of Canada and the Department of Finance in late February 1991, and initially stipulated a downward path for inflation from a rate of roughly six percent to two percent by 1995.[24] A target or target path beyond 1995 was not specified beyond a general statement about further reductions to achieve price stability, thought to be less than two percent.[25] By December 1993, inflation had fallen to two percent and the Bank of Canada and the Government of Canada agreed to a framework going forward of a two percent mid-point target within a 1–3 percent range.[26] This framework was to be in place until 1998 when it would be reviewed; subsequently, this framework has been renewed since numerous times with the same target and range.

The fact that the initial phase of the inflation targeting framework was not a two percent mid-point target with a 1–3 percent range, as it is currently described, means the start date of the current inflation targeting framework is properly after the initial announcement in 1991. Rather, the natural starting date is 1994M1,

[24] See Bank of Canada (1991).

[25] In Bank of Canada (1991), the Bank referenced existing analysis that supported this definition of price stability but indicated further work is still required.

[26] Bank of Canada (1993).

immediately following the first framework agreement with the two percent midpoint target specified. On this basis, we set the start of the inflation targeting period to be 1994M1.

The Bank of Canada's inflation target has always been defined in terms of the 12-month rate of change in CPI All-items (hereafter, simply CPI) and this measure is the focus of our empirical work. The Bank has also referenced various *core* measures of inflation over the years as a guide to underlying inflation. These measures are meant to remove the more volatile relative price changes that can affect total CPI inflation in order to provide a better guide for monetary policy.[27] Given the important role core inflation has had in Canadian monetary policy, we also examine the behaviour and implied policy horizons for one of these core measures, CPIX. While CPIX is not currently the preferred measure of core inflation, we view it as a compromise choice in that it has been part of the discussions around core inflation measures since 2001 and is quite closely related to the measure of core used prior to 2001.[28]

We measure inflation as year-on-year log differences, which is an approximation to the published rate of inflation targeted by the Bank of Canada and the CPIX measure of inflation. That is,

$$\pi_t = \ln P_t - \ln P_{t-12}$$

where P_t is the CPI or CPIX price index. Both series are seasonally unadjusted.

The methods here are based on the assumption that the underlying economic environment and behaviour of the central bank are stable and give rise to a process for inflation, as in equation (10). For this reason, we have concerns about using the full sample of data available to us (at the time of our analysis, through to 2022M9). The onset of the pandemic led to a rapid disinflation and a subsequent large global and local inflation shocks. This is evident in Figure 1, which presents CPI and CPIX inflation through to 2022M9, and sees CPI inflation peak well in excess of 7 percent. To address these concerns, we limit our analysis to the sample 1994M1–2019M12, excluding the pandemic and recent events. Thus, the full sample is 1994M1–2019M12, a total of $T = 312$ observations.

To further support this choice, we use the Bai and Perron (1998) test for a single unknown structural break in mean inflation (for both CPI and CPIX) over the sample 1994M1–2022M9 (342 observations). The trim was set to ten percent, which leaves just over thirty observations for the smallest sub-samples. For both inflation series, the estimated break point is 2019M11; for CPI inflation, the supW test-statistics is 50.72; for CPIX inflation, 88.09. Both are well in excess of the one-percent critical value of 13.00. As there is uncertainty around the estimate of the breakpoint, it seems reasonable and neater to end the estimation sample in 2019M12.[29]

With the full sample determined as 1994M1–2019M12, we further consider possible instability in the economic environment or central bank behaviour by considering three sub-samples based upon the tenures of three past governors of the Bank of Canada: Dodge, Carney, and Poloz; these samples are 2001M2–2008M1,

[27]See Macklem (2001) and Khan et al. (2015) for early and later discussions on core measures of inflation and their role in monetary policy for the Bank of Canada. Khan et al. (2015) provides a summary of the preferred measures of core inflation prior to the pandemic. See Bank of Canada (2022) for the Bank's recent assessment of the performance of different core measures.

[28]See Macklem (2001).

[29]The differences in estimated means pre- and post-2019M12 are considerable. For CPI inflation, the mean prior to the break is 1.78 percent; post breakpoint, 3.28 percent. For CPIX inflation 1.74 and 2.94 respectively. Of course, we are currently mid-shock and in due course we can expect the mean to return to somewhere close to two percent, consistent with the Bank of Canada's inflation target. For empirical purposes, however, it seems best to leave the current episode out of the analysis.

2008M2–2013M5, and 2013M6–2020M6.[30] We see this as a natural means of identifying sub-samples, as they allow for possible changes in overall policy. Any differences between the sub-samples, however, cannot be identified as solely due to the change in leadership. Changes in the nature of the shocks as well as the overall economic environment can also influence the results. Our methods does not allow us to separately identify these effects.

For the dynamic models that we estimate, different lag lengths p are required. For each of the samples of data used for estimation, we let $t = p-1, \ldots, 0, 1, \ldots T$, a total of $T+p$ observations with $t = 1$ corresponding to the first observation of the full sample (1994M1) or the various sub-samples.

Summary statistics for the full sample, sub-samples, and the source information for the two inflation measures are provided in Table 1.

The underlying price indices, CPI and CPIX, exhibit considerable seasonality that is, to some extent, controlled for by the seasonal differencing used to construct the inflation measure. The resultant inflation rate series, though, continue to exhibit a great deal of persistence at seasonal and near-seasonal frequencies, as evident from the autocorrelation functions in Figure 3 and the partial autocorrelation functions in Figure 4. The series also exhibit considerable non-normality, as evident from the Shapiro and Francia (1972) test reported in Table 1 and Figure 5, which shows the q-q plots for both series. These features of the data bear on the appropriate time series models for inflation.

4.2 Time series models

For each inflation series, we need time series models for the conditional forecasts that underlie our policy horizon estimates. Our objective is to find a parsimonious model that fits the dynamics of the series as best as possible. Inspection of the autocorrelation and partial autocorrelation functions in Figures 3 and 4 provide some guidance on the appropriate models.

For both inflation series, the autocorrelation functions decay in a manner consistent with an autoregressive model. Inspection of the partial autocorrelation functions suggests, though, that the autoregressive process has a large number of lags. Indeed, standard tests of lag parameter significance suggest lag lengths in excess of thirty-six — consistent with a visual inspection of the partial autocorrelation functions. To explore this, we consider AR models of lag lengths 12, 24, and 36 and examine the residuals for serial correlation. Table 2 reports the results from the serial correlation tests of order 1, 12, and 24. In all cases, there is strong evidence of lower- and higher-order serial correlation in the residuals.

Since simple AR models do not adequately model monthly inflation and since neither the autocorrelation or partial autocorrelations are consistent with simple MA models, we next turn to ARMA models. Further, because of the seasonal or near-seasonal spikes in partial autocorrelations, we also consider multiplicative seasonal ARMA models. The general multiplicative seasonal ARMA model we start with is:[31]

$$\rho(L^p)\rho_{12}(L^P)(\pi_t - \pi_0) = \theta(L^q)\theta_{12}(L^Q)\varepsilon_t$$

[30] Governor Poloz served from 2013M6–2020M6; however, as with the full sample, we limit the sample to end in 2019M12.

[31] By construction, our model imposes a seasonal difference on the log price level, so strictly this is a seasonal ARIMA model. See Harvey (1981) for a detailed discussion of these models. We also consider additive seasonal AR and MA components, that is ARMA lag lengths of 12, but the multiplicative models seem to work better. We recognize that there may be other suitable time series specifications for these series.

where

$$\rho(L^p) = (1 - \rho_1 L - \dots \rho_p L^p) \qquad \rho_{12}(L^P) = (1 - \rho_{12,1}L^{12} - \dots - \rho_{12,P}L^{P\cdot 12})$$
$$\theta(L^q) = (1 - \theta_1 L - \dots \theta_q L^q) \qquad \theta_{12}(L^Q) = (1 - \theta_{12,1}L^{12} - \dots - \theta_{12,Q}L^{Q\cdot 12})$$

We begin by examining models with $P = Q = 1$ and $p = q = 2$ and refine each model (inflation measure and sub-sample) based upon the statistical significance of the individual parameters.[32] A key concern is to model fully the dynamics as best as possible; to this end, we examine tests for white noise of the residuals at $6, 12,$ and 24 lags as part of the model selection process. In one instance — the full sample model for CPI inflation — the presence of serial correlation leads to extending the seasonal models to $P = Q = 2$ and selecting a model with $P = 2$.

The results for CPI inflation for the four different samples are presented in Table 3. The results for CPIX inflation are reported in Table 4. For each sample, the table reports the estimated coefficients for the selected model and the associated P-value. In all instances, the coefficients are statistically significant at the five percent test level. The table also reports 95 percent confidence intervals for the estimated mean inflation rate, which will provide some guidance to the analysis below. Also reported for each model are Q-tests for white noise and the associated P-value.[33]

For CPI inflation, the full sample model is an ARMA(1,1) with a multiplicative seasonal MA(1) term. The estimated mean is 1.8 percent, with a 95 percent confidence interval of 1.7 to 2.0 percent. The estimated mean and range are slightly below the Bank of Canada's inflation target. The standard deviation of the residuals is 0.30 percent, suggesting a reasonable fit for the model. A comparison of the fitted and actual CPI inflation is presented in Figure 6 that further supports this conclusion. The tests for white noise residuals at $m = 6, 12,$ and 24 show no evidence of serial correlation, providing further evidence that the model adequately captures CPI inflation dynamics. There are two notable features of the estimated dynamics. The first is the significant amount of persistence, with an AR(1) coefficient of 0.96. The second is the reasonably large seasonal MA terms of -0.87. Both of these will contribute to the estimated policy horizons, in particular the estimated persistence.

For the sub-sample models for CPI inflation, there are points of similarity as well as differences, relative to the full sample, evident from the estimated coefficients and residual diagnostics. To organize the discussion, it is useful to first note that the two later sub-samples, 2008–2013 and 2013–2019, are very similar, with the only points of departure the amount of persistence (slightly less in the latter sample) and the standard deviation of the residuals (again, slightly less in the latter sample). As well, both sub-sample models show no evidence of serial correlation in the residuals. The two later sub-samples are also broadly comparable to the full sample model, the main distinctions being that the sub-sample models have lower mean inflation rates and have much less persistent dynamics.

In contrast, however, the first sub-sample covering 2001–2008 differs substantially from the full sample (and subsequent sub-samples) in a number of ways. First, the estimated mean is much higher at 2.2 percent (and confidence interval of 2.0–2.4 percent). Second, while persistence governed by the AR(1) term is comparable to the full sample, the two seasonal multiplicative AR terms suggest quite different dynamics, as do the

[32] For some of the samples, this estimation of this general model failed and we moved to more parsimonious models.

[33] Although not reported, there is some evidence of non-normality of the residuals for all of these models, which qualifies any inferences about the models.

MA coefficients. Perhaps the most troubling difference, however, is the strong evidence of residual serial correlation at $m = 6$ and $m = 12$.[34]

The differences between the sub-sample models and the full-sample model, as well as the differences across sub-sample models, suggest that parameter instability for the full-sample model may be a concern and qualify our analysis. The sub-samples are, however, relatively small, which itself qualifies the sub-sample estimation. Further, as we will see below, despite the quite different dynamics, the policy horizons from the 2001–2008 sub-sample while different, are still broadly comparable to the full sample estimates.

For CPIX inflation, reported in Table 4, there are again similarities and differences across the CPIX inflation models (full sample and sub-samples) as well as with respect to CPI inflation. As our principal focus is CPI inflation, we limit the details of our discussion of the CPIX models to some general conclusions. For the full sample model, the overall structure of the model is relatively similar to the CPI inflation model. There is evidence of considerable persistence, similar to CPI inflation, as well as a similarly-sized seasonal MA term. The principal difference in structure is the presence of a multiplicative seasonal AR term. The overall fit is also similar, though with a lower mean inflation rate and standard deviation of the residuals, which is to be expected given that the CPIX inflation is considerably less volatile than CPI inflation. A comparison of actual and fitted CPIX inflation is presented in Figure 7.

For the sub-sample models of CPIX inflation, we observe more heterogeneity in model structures for the sub-samples as well as with respect to the full-sample model. There is considerable variation in persistence in terms of the AR and seasonal AR terms and the MA structures for the sub-samples are notably different from the full sample model. Finally, we again see evidence of serial correlation in the 2001–2008 sub-sample, suggesting that as with CPI inflation, this sub-sample is difficult to model adequately. Overall, it appears that finding a stable model of CPIX inflation over this sample is difficult, which does qualify the policy horizon estimates presented below. As with CPI inflation, comparison of the policy horizon results for the different sub-samples should allow us to determine whether the different estimated models materially affect the policy horizon measures.

The final task before we examine the distribution of policy horizons for each estimated model is to select a relevant range of tolerances for convergence to target. As noted above, models with autoregressive components will only converge to the unconditional mean at an infinite horizon, so some finite tolerance criteria is required. At a minimum, the tolerance criteria should reflect the uncertainty associated with the unconditional mean estimate of π_0, the point to which the forecasts converge eventually. We use the 95 percent confidence intervals for π_0, reported in Tables 3 and 4 to guide our choice for the smallest tolerance level to consider. For CPI inflation, the confidence intervals range from $\pm$ 0.15 percent (full sample) to 0.39 percent (2008–2013 sample). From this, we select 0.30 as our smallest tolerance level and construct a set of tolerance ranges of $\{0.30,\ 0.50,\ 0.70\}$ for CPI inflation. As we shall see, these ranges provide policy horizons comparable to those of the Bank of Canada. For CPIX inflation, the confidence intervals range are from $\pm$ 0.08 percent (full sample) to $\pm$ 0.28 percent (2001–2008 sample). Based on these, we use $\{0.10,\ 0.20,\ 0.30\}$ as the set of tolerance ranges for CPIX inflation.

[34] We consider a number of alternative models but are unable to model resolve this problem.

4.3 Estimated Policy Horizons

4.3.1 Full Sample Results

We now use the full sample models from Tables 3 and 4 to construct estimates of the policy horizons based on equation (14), amended to use the unconditional mean of the inflation model, π_0:

$$|E_t\pi_{t+j} - \pi_0| < \mu \qquad \forall j \geq N_t(\chi_t; \rho, \Theta, \pi_0; \mu), \quad t = 1 \ldots T \tag{15}$$

We consider values for μ, the tolerance level, from the two sets identified above (depending upon the inflation measure used) to identify the policy horizon N_t.

The procedure for estimating the policy horizons is as follows. We use the model of interest to construct a rolling set of forecasts for $j = 1 \ldots 72$ months ahead, starting with the first observation of the estimation sample, $t = 1$ (which is 1994M1 for the full sample), and prior lags (as determined by the model) as the conditional information set, χ_t. We then move through the full estimation sample to period T, constructing conditional forecasts $E_t\pi_{t+j},\ j = 1 \ldots 72$ for each $t = 1...T$.[35]

For each estimated conditional forecast, we use equation (15) to construct an estimated policy horizon $\hat{N}_t(\chi_t; \hat{\rho}, \hat{\Theta}, \hat{\pi}_0; \mu),\ t = 1 \ldots T$. Written in this manner, it is clear that the only source of variation in $\hat{N}_t$ over the sample, for a given estimated model, is the conditional information set χ_t. That is, the variation in policy horizons that we will observe are functions only of the recent history of inflation.

For each estimated model, the procedure above gives a set of estimated policy horizons $\hat{N}_t$, $t = 1, \ldots T$. We could limit our discussion to the distribution of the entire set of estimated policy horizons; however, in many instances when inflation is already close to target, the policy horizons will be very short (as little as one month), indicating that projected inflation is within the tolerance band immediately and at all subsequent horizons. As an alternative, we isolate instances where monetary policy has work to do; that is, when current inflation is outside of the Bank of Canada's 1–3 percent control range. This is the set of policy horizons defined as $\{\hat{N}_t\ |\pi_t \notin [1.0, 3.0]\}$. Below we examine the distribution for the full set and the restricted set to get a better sense of how policy horizons depend upon the inflation situation at the time.

Table 5 reports the distribution statistics for the full sample policy horizons for CPI and CPIX inflation; the top panel for the full set and the bottom panel for the restricted set. The full sample set consists of 312 policy horizon observations and the restricted set consists of 72 policy horizon estimates, which is 23 percent of the full sample set. Figures 8 and 9 report the histograms for the full set of policy horizons for different values of μ for CPI and CPIX inflation. Figures 10 and 11 do the same but for the restricted set of policy horizons.

We focus our discussion on CPI inflation, which is the basis for the Bank of Canada's inflation target. Recall from the earlier discussion that the Bank views the typical policy horizon to be six to eight quarters, or 18–24 months, which provides a benchmark for interpreting our results. We can also compare our policy horizon distributions with those of the optimal policy horizon studies for Canada, Cayen et al. (2006) and Basant-Roi and Mendes (2007). For simplicity, we focus on the summary of the results for these two studies reported in Coletti et al. (2006): policy horizon means of 18 and 21 months, respectively, and ranges of

[35] The choice of 72 for the forecast horizon ensures that, for our tolerance ranges and estimated models, we are able to identify a policy horizon; that is, our estimated policy horizons do not exceed 72.

$[12, 33]$ months and $[6, 27]$ months, respectively.[36] These ranges are constructed from 90 percent confidence bands from the empirical distribution of simulated optimal policy horizons.

We start with the complete set of policy horizons reported in the upper panel of Table 5. The mean policy horizon lengths for CPI inflation are 19, 25, and 35 months for $\mu = 0.70$, 0.50, and 0.30 respectively. The first two mean policy horizons are consistent with the Bank's stated typical policy horizon of 18–24 months, as well as the simulation studies of optimal policy horizons.[37] The distributions of the policy horizons, however, are not symmetric nor are they unimodal and for all tolerance levels, the mean and median differ substantially, with the median exceeding the mean.[38] These features are evident in the histograms in Figure 8 and the statistics reported in Table 5. Given these features of the distribution, the median policy horizons are a somewhat more useful measure of central tendency than the mean and these are 22, 30, and 42 months for $\mu = 0.70$, 0.50, and 0.30 respectively.

These results, as expected, depend upon the choice of μ and we can approach this in two ways. First, we can decide that a particular value of μ is most relevant and, with this choice, identify the *de facto* policy horizon for the Bank of Canada associated with our models of inflation. For this purpose, we select $\mu = 0.50$. This is a larger value than implied by the uncertainty surrounding the unconditional mean estimates for inflation; however, it allows, in an arbitrary way, for other modelling uncertainty, including parameter instability. For $\mu = 0.50$, our *de facto* median policy horizon for CPI inflation is 30 months, two and one-half years. This is significantly longer than the Bank's typical policy horizon of 18–24 months or one and one-half to two years.

The alternative approach, following the discussion in Mu (2022), is to select μ to align the estimated median policy horizon with the Bank's typical horizon. The tolerance level $\mu = 0.70$ has a median policy horizon of 22 months, which matches the Bank's typical horizon quite well. In other words, the Bank's policy horizon of 18–24 months refers to the length of time to return expected inflation to target, plus or minus 0.70 percent. This is quite a wide tolerance band but is one means of characterizing the uncertainty that underlies the Bank's policy statements and projections they report over the policy horizon.[39]

While the median policy horizon provides some information, the range and variation of the estimated policy horizons are also of interest. Table 5 reports the standard deviation and coefficient of variation for the set of policy horizons but given the nature of the empirical distributions, it is more useful to consider the range between the 5^{th} and 95^{th} percentiles. This measure is comparable to the reported ranges in the studies cited above, as well as the histograms reported in Figure 8. For $\mu = 0.50$, which we will focus on, the range for these percentiles is [1, 53]. This is significantly wider than the ranges reported in Coletti et al. (2006). This also puts in perspective the Bank of Canada's typical policy horizon, with our upper band estimate of 53 months more than double the typical 18–24 month horizon.

To this latter point, Figure 8 provides a starker perspective. There is very little mass of the distribution

[36] These studies use quarterly data and report quarterly policy horizons that we have expressed in months. Basant-Roi and Mendes (2007) report policy horizons with and without the presence of housing-price bubbles. We are using the latter for comparison.

[37] The tolerance level used Cayen et al. (2006) and Basant-Roi and Mendes (2007) is 0.10, smaller than we use here. These studies use theoretical models, which likely generate smoother and less persistent inflation paths , allowing for a lower tolerance level. Further, these studies use a quarterly inflation measure that may also influence the choice of μ. Finally, these studies are examining impulse response functions, subject to multiple shocks, rather than conditional forecasts. The critical difference being the role of the conditional information, χ_t.

[38] This relationship is reversed for CPIX inflation, with the median less than the mean.

[39] In recent years, the Bank of Canada does not provide uncertainty measures associated with their inflation projections at the time of policy decisions. See Mu (2022) for a discussion of how the presentation of inflation projections have evolved since the early 2000s.

centred on 18–24 months. Rather, there is a significant proportion of policy horizons of roughly one year or less, associated for the most part with inflation rates close to target, and a large proportion of policy horizons centred around (just below) 36 months. From these results, one might reason that the typical policy horizon is conditional and better described as follows. When current and recent inflation conditions are close to target, the typical policy horizon is one year or less. When recent conditions have inflation well away from target, the typical policy horizon is three years.

We further explore this aspect in the bottom panel of Table 5, which reports the restricted set of policy horizon measures, those for which π_t — the last element in χ_t — is outside of the inflation control range of 1–3 percent. In this case, for $\mu = 0.50$, the median policy horizon for CPI inflation is 34 months, with a range of six to 53 months. As before, we get a better sense of the distribution from the histogram, reported in Figure 10. We see that while there is a large proportion of policy horizons centred on 36 months, we also see a significant proportion that are shorter in length, around 12 months, as well as beyond at 48 months. In summary, the nature of the distribution of our policy horizon estimates is such that it is difficult to accurately characterize a *typical* policy horizon and that focusing on some measure of central tendency obscures the variation and dependence on current conditions. Moreover, policy horizons depend not just on current inflation but on the recent history of inflation.

Table 5 also reports the policy horizon summary statistics for CPIX inflation, for both the full set as well as the restricted set. The histograms are in Figures 9 and 11. As expected, these policy horizons are considerably shorter. The median policy horizons range from six months ($\mu = 0.30$) to 23 months ($\mu = 0.10$). The associated ranges are still quite wide, though. For $\mu = 0.10$, for example, the range (5th and 95th percentile) for the full set of policy horizons range is 10 months to 52 months. The dispersion is also evident in the histograms. A further distinguishing feature, relative to CPI inflation, is that the median is now consistently less than the mean policy horizon, indicating the majority of observations are massed on shorter horizons but with a long righthand tail to the distribution.

4.3.2 Sub-sample Results

We now consider sub-samples associated with three recent Governors of the Bank of Canada: Governors Dodge, Carney, and Poloz. To focus and limit the discussion, we only consider the policy horizons for CPI inflation and report only those for the mid-point tolerance level used for the full-sample, $\mu = 0.50$. The models used to construct the forecasts are those reported in Table 3. As with the full-sample results, we report summary statistics for the full set of policy horizons as well as for the restricted set, where the last inflation observation of the information set χ_t is strictly outside of the 1–3 percent inflation control range.

The summary statistics for each sub-sample are reported in Table 6 and the histograms are reported in Figure 12. From the table, we see that the median (or mean) policy horizons are short relative to those from the full-sample estimation. If we focus on the restricted set of policy horizons, the longest median policy horizon is for the Dodge sub-sample at twenty months. For Carney, it is nine months, and for Poloz five months.

Between the three sub-samples themselves, there is considerable variation in the distributions of the policy horizons. This is evident in the different mean or median lengths of the policy horizons as well as in the range of policy horizons. The Carney and Poloz sub-samples have a much narrower distribution than the

Dodge sub-sample and the full-sample results. This is also evident in the histograms reported in Figure 12.

The shorter policy horizons for the Poloz sub-sample, and to a lesser extent the Carney sub-sample, are worth considering further. The Poloz sub-sample differs from all of the other sub-samples in terms of the low values for the mean and variability of CPI inflation, as seen in Table 1. This is also evident in Table 6, where we see that there are just five observations out of the 79 observations in the sample that are outside of the 1–3 percent control range; this is just six percent of the observations. For the full sample, recall the number is 23 percent; for Dodge, 18 percent; and for Carney, 39 percent. This relatively stable period of inflation is an important factor in the associated short policy horizons.

The Carney sub-sample, though, is perhaps the most interesting of the three. Although it is characterized by the highest percentage of inflation observations outside of the 1–3 percent inflation control range — by a factor of roughly two relative to the full-sample and the Dodge sub-sample — the median policy horizon is much shorter at nine months relative to 34 months for the full sample and 20 months for the Dodge sample.

There are two perspectives to understand this result, though they are necessarily closely related. First, the difference in policy horizons may be understood by the differences in the estimated time series models. From this perspective, we see from Table 3 that the notable distinguishing feature for the Carney model (and the Poloz model) is much less persistence evident in the parameter estimates relative to the full sample and the Dodge sample models. Second, the differences in policy horizons may be understood by the differences in the conditional information sets, χ_t, across the samples. In both cases, the differences may reflect the distribution of the underlying shocks to the economy, the structure of the economy, and the behaviour of the central bank. Our methods, however, cannot separately identify these influences.

As a means of examining how the different estimated model contributes to the lower horizons from the Carney model, we perform the following thought experiment. We use this model to measure policy horizons over the full sample of data (using the tolerance level of $\mu = 0.50$) and compare it to the policy horizons from the full sample model. For the 72 observations of the restricted set of policy horizons, the median policy horizon using the Carney model for inflation is 12 months, with a range (5^{th} and 95^{th} percentile) of 3–21 months. This is substantially different from the median for the full sample of 34 months with a range of 6–53 months.

An important factor in these different results is the lower persistence in the model estimated over the Carney sample and how this interacts with the patterns of inflation over the full-sample. To get a sense of this contribution, Figure 13 reports a selection of conditional forecasts that illustrate the differences in convergence from the two models. The much quicker convergence to target from the Carney sample model is particularly evident in the first panel, which shows the conditional forecast for 1995M1 and subsequent periods. But an important feature of these conditional forecasts, which is not immediately evident from the estimated models themselves, is that the conditional information sets matter. For the conditional forecasts in the top panel of Figure 13, the higher persistence of the full sample model is very evident; for those in the bottom panel, much less so. This role for the conditional information for the policy horizon estimates is an important distinguishing feature of our approach compared to studies discussed earlier that use impulse response functions.

The principal conclusion from these sub-sample results is that the variation in the estimated models (coefficients, residual variance) is sufficient to give rise to quite different distributions of policy horizons. From an empirical perspective, this means that to estimate *de facto* policy horizons, it is important to consider the stability of the estimated model. We have done so here by arbitrarily selecting three sub-sample periods to examine the issue and clearly these results qualify our policy horizon estimates. Future work might focus more explicitly on testing for, or modelling, parameter instability. From a practical perspective, our results suggest that policy horizons likely vary over time, not just because of the paths of inflation that inform the conditional forecasts but quite likely due to changing behaviour of the central bank and the underlying macroeconomic structure and forces for inflation.

5 Conclusion

We build on existing literature to estimate policy horizons for monetary policy. Our results contrast from previous work by using conditional forecasts to estimate a set of policy horizons over the sample, rather than rely on impulse response functions. We further differ by estimating *de facto* policy horizons rather than characterizing optimal policy horizons, either through estimated models or calibration.

We examine a number of different samples and consider both CPI and CPIX inflation. The representative result, one that is most comparable to the typical policy horizon of the Bank of Canada, is a median policy horizon of 30 months, considerably longer than the typical policy horizon of 18–24 months used by the Bank of Canada.

In developing our methods and estimating the policy horizons, we explore a number of issues that provide a deeper understanding of monetary policy horizons. The first is that the horizons are quite sensitive to the choice of tolerance range for convergence to target, which is not too surprising. Relatedly, though, is that the tolerance range needed to get estimated policy horizons with a median roughly consistent with the Bank of Canada's typical policy horizon, 18–24 months, is quite large, $\pm$ 0.70 percent.

The second and more important issue is that our results call in to question the usefulness of identifying a typical policy horizon; or, put differently, to focus on either the mean or median policy horizon from an estimated set. This is relevant for our approach but also the simulation approach used by Cayen et al. (2006) and Basant-Roi and Mendes (2007). The distributions for the sets of policy horizons we measure are not unimodal and are heavily influenced by current conditions (i.e., the conditional information set). What we observe, for a given estimated model and sample, is a clustering of policy horizon estimates of short duration (e.g. one year), and a clustering of policy horizon estimates at much longer horizons (e.g. three years).

In effect, this is one form of flexibility that central banks such as the Bank of Canada refer to and is indeed evident in their communications, as we noted in the introduction — specifically, that the time needed to return inflation to target will vary depending upon the nature and persistence of the shocks. What our results add to this, though, are measures of the time that this might in fact often involve (three years or longer), and, as noted, whether means/medians are useful measures to summarize the length of time expected to return to target. This matters for communications purposes as it raises questions about the credibility and meaning of policy horizons. It also provides support to the idea that central banks should provide detailed conditional forecasts, including the associated uncertainty with the forecasts.

Our methods are simple and our treatment of flexibility is limited. Although this may qualify our estimates of policy horizons, we expect that richer empirical environments, for example multivariate models, would likely deliver similar conclusions, particularly about the distributions of the policy horizons. This is conjecture though and requires further research. The other direction for further research is to better explore the stability of the models for inflation, including modelling inflation using time series models with time-varying parameters. The quite distinctive results we get with the small sub-samples we investigate suggests that this may indeed be a significant qualification to our results. Relatedly, a more ambitious direction would be to develop a structural model that embeds concepts of flexibility directly into the inflation targeting framework and the associated policy horizons.

References

Bai, J., and P. Perron (1998) "Estimating and testing linear models with multiple structural changes," *Econometrica* 66(1), 47–78

Bank of Canada (1991) "Press release: Targets for reducing inflation," *Bank of Canada Review* pp. 5–6

——— (1993) "Statement of the Government of Canada and the Bank of Canada on monetary policy objectives," Press Release 22 December 1993, Bank of Canada

——— (2022) *Monetary Policy Report October 2022*, Bank of Canada

Basant-Roi, M., and R. R. Mendes (2007) "Should central banks adjust their target horizons in response to house-price bubbles?," Discussion Paper 2007-4, Bank of Canada

Batini, N., A. Justiniano, P. Levine, and J. Pearlman (2006) "Robust inflation-forecast-based rules to shield against indeterminacy," *Journal of Economic Dynamics and Control* 30(9), 1491–1526

Batini, N., and E. Nelson (2001a) "The lag from monetary policy actions to inflation: Friedman revisited," *International Finance* 4(3), 381–400

——— (2001b) "Optimal horizons for inflation targeting," *Journal of Economic Dynamics and Control* 25(6), 891–910

Bernanke, B. S., T. Laubach, F. S. Mishkin, and A. S. Posen (1999) *Inflation Targeting: Lessons from the International Experience*, Princenton, NJ: Princeton University Press

Cayen, J.-P., A. Corbett, and P. Perrier (2006) "An optimized monetary policy rule for totem," Working Paper 2006–41, Bank of Canada

Coletti, D., J. Selody, and C. Wilkins (2006) "Another look at the inflation-target horizon," *Bank of Canada Review*

Davig, T., and A. Foerster (2021) "Communicating monetary policy rules," Working Paper 2021-12, Federal Reserve Bank of San Francisco

Friedman, M. (1961) "The lag in effect of monetary policy," *Journal of Political Economy* 69(5), 447–466

——— (1972) "Have monetary policies failed?," *The American Economic Review* 62(1/2), 11–18

Gabriel, V. J., P. Levine, and C. Spencer (2009) "How forward-looking is the Fed? Direct estimates from a 'Calvo-type' rule," *Economics Letters* 104(2), 92–95

Giannoni, M. P., and M. Woodford (2003) "How Forward-Looking is Optimal Monetary Policy?," *Journal of Money, Credit, and Banking* 35(6, Part 2), 1425–1469

Harvey, A. C. (1981) *Time Series Models*, Oxford, UK: Philip Allan

Hunt, J. (2022) "Remit for the Monetary Policy Committee November 2022," Letter from the Chancellor to the Governor and Attachment, HM Treasury, United Kingdom (available at: https://www.bankofengland.co.uk/letter/2022/november/mpc-remit-2022)

Khan, M., L. Morel, and P. Sabourin (2015) “A comprehensive evaluation of measures of core inflation for Canada,” Discussion Paper 2015-12, Bank of Canada

King, M. (1994) “Monetary policy in the UK,” *Fiscal Studies* 15(3), 109–128

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

Macklem, T. (2001) “A new measure of core inflation,” *Bank of Canada Review* pp. 3–12

Mu, G. (2022) *Inflation Targeting and Policy Horizons*, MA Thesis, University of Victoria

Otto, G. D., and G. M. Voss (2011) “What do the RBA’s forecasts imply about its preferences over inflation and output volatility?,” *Economic Record* 87(279), 509–524

——— (2014) “Flexible inflation forecast targeting: Evidence from Canada,” *Canadian Journal of Economics/Revue canadienne d’économique* 47(2), 398–421

Rowe, N., and J. Yetman (2002) “Identifying a policymaker’s target: an application to the Bank of Canada,” *Canadian Journal of Economics/Revue canadienne d’économique* 35(2), 239–256

Royston, J. P. (1983) “A Simple Method for Evaluating the Shapiro-Francia W′ Test of Non-Normality,” *Journal of the Royal Statistical Society. Series D (The Statistician)* 32(3), 297–300

Shapiro, S. S., and R. S. Francia (1972) “An approximate analysis of variance test for normality,” *Journal of the American Statistical Association* 67(337), 215–216

Smith, G. W. (2019) “Two Types of (Slight) Flexibility in Bank of Canada Projections, 2003–2019,” *Canadian Public Policy/Analyse de Politiques* 45(3), 366–376

Svensson, L. E. (1999) “Inflation targeting as a monetary policy rule,” *Journal of Monetary Economics* 43(3), 607–654

Woodford, M. (2007) “Forecast Targeting as a Monetary Policy Strategy: Policy Rules in Practice,” Working Paper 13716, National Bureau of Economic Research, Cambridge, MA

Table 1: Summary Statistics

	Obs.	Mean	Median	SD	Skewness	Kurtosis	z	P-value
1994M1–2019M12								
CPI Inflation	312	1.782	1.833	0.864	-0.025	3.786	2.000	0.023
CPIX Inflation	312	1.736	1.713	0.414	0.283	3.165	1.714	0.043
2001M2–2008M1								
CPI Inflation	84	2.216	2.167	0.807	0.625	4.077	2.824	0.002
CPIX Inflation	84	1.958	1.875	0.422	0.656	3.251	2.194	0.014
2008M2–2013M5								
CPI Inflation	64	1.672	1.708	1.078	-0.334	2.829	0.645	0.259
CPIX Inflation	64	1.661	1.678	0.332	-0.044	2.794	-2.185	0.986
2013M6–2019M12								
CPI Inflation	79	1.640	1.541	0.499	0.267	2.333	1.342	0.090
CPIX Inflation	79	1.654	1.655	0.411	-0.259	2.107	1.476	0.070

Variable definitions and source:

(1) CPI All-items (2002=100). Statistics Canada Table 18100004 [v41690973]

(2) CPIX, All-items excluding eight of the most volatile components as defined by the Bank of Canada and excluding the effect of changes in indirect taxes. Statistics Canada Table 18100256 [v112593702].

Series are inflation rates, measured as year-on-year percent change (log).

Notation: Obs. is the number of observations; SD the standard deviation; and z is the standard normal test statistic for the Shapiro and Francia (1972) test for normality based on Royston (1983). The P-value is reported in the adjacent column.

Table 2: Autoregressive Models – Tests for Serial Correlation

Sample: 1994M1–2019M12	Order of serial correlation		
	$m = 1$	$m = 12$	$m = 24.$
CPI Inflation			
AR(12)	0.0000	0.0000	0.5623
AR(24)	0.0000	0.0001	0.0000
AR(36)	0.0000	0.0207	0.0000
CPIX Inflation			
AR(12)	0.0000	0.0000	0.0017
AR(24)	0.0000	0.0000	0.0000
AR(36)	0.0044	0.0021	0.1505

Numbers are P-values for an LM test statistic for serial correlation in the residuals of order m; the test statistics is $\chi^2(1)$.

Table 3: CPI Inflation Models

	Sample			
	1994M1–2019M12	2001M2–2008M1	2008M2–2013M5	2013M6–2019M12
π_0	1.8221	2.2105	1.6592	1.6478
	0.0000	0.0000	0.0000	0.0000
95% C.I.	[1.67, 1.98]	[2.04, 2.38]	[1.27, 2.05]	[1.45, 1.85]
ρ_1	0.9583	0.9670	0.8959	0.8559
	0.0000	0.0000	0.0000	0.0000
θ_1	0.1116			
	0.0465			
ρ_2		-0.2091		
		0.0501		
$\theta_{12,1}$	-0.8737		-0.6954	-0.5554
	0.0000		0.0019	0.0000
$\rho_{12,1}$		-0.8187		
		0.0000		
$\rho_{12,2}$		-0.4400		
		0.0000		
σ	0.3044	0.3512	0.3722	0.2618
$Q(6)$	3.0419	6.7250	5.0480	2.6499
	0.2185	0.0095	0.1683	0.4488
$Q(12)$	7.6906	14.0425	7.5311	4.5097
	0.4643	0.0504	0.5820	0.8748
$Q(24)$	11.8060	22.5088	21.8043	22.4105
	0.9226	0.2597	0.4109	0.3762
Obs.	312	84	64	79

See text for coefficient definitions; σ is the standard deviation of the residual. Numbers below coefficient estimates are P-values. For π_0 estimates, 95 percent confidence intervals are reported. $Q(m)$ is the Q test for white noise with m lags, distributed $\chi^2(m-k)$ where k is the number of estimated coefficients; see Harvey (1981).

Table 4: CPIX Inflation Models

	Sample			
	1994M1–2019M12	2001M2–2008M1	2008M2–2013M5	2013M6–2019M12
π_0	1.7365	1.9180	1.6684	1.6413
	0.0000	0.0000	0.0000	0.0000
95% C.I.	[1.66, 1.81]	[1.64, 2.20]	[1.57, 1.77]	[1.38, 1.90]
ρ_1	0.9319	0.8642	0.6227	0.9237
	0.0000	0.0000	0.0000	0.0000
θ_1	-0.1402			
	0.0345			
$\rho_{12,1}$	0.1992	-0.3767	-0.5046	
	0.0160	0.0000	0.0000	
$\theta_{12,1}$	-0.7809			-0.3947
	0.0000			0.0002
σ	0.1764	0.2156	0.2137	0.1483
$Q(6)$	1.0696	2.3814	1.9645	4.8899
	0.3010	0.4971	0.5798	0.1800
$Q(12)$	8.0749	7.7523	9.8408	10.6291
	0.3260	0.5593	0.3635	0.3020
$Q(24)$	16.0898	35.0061	25.7524	26.5543
	0.6513	0.0282	0.2160	0.1861
Obs.	312	84	64	79

See text for coefficient definitions; σ is the standard deviation of the residual. Numbers below coefficient estimates are P-values. For π_0 estimates, 95 percent confidence intervals are reported. $Q(m)$ is the Q test for white noise with m lags, distributed $\chi^2(m-k)$ where k is the number of estimated coefficients; see Harvey (1981).

Table 5: Policy Horizons — Full Sample

	Mean	Median	S.D.	C.V.	p5	p95	No. PH
All Policy Horizons							
CPI Inflation							
$\mu = 0.30$	34.53	42.00	19.02	0.55	8	65	312
$\mu = 0.50$	24.84	30.00	16.21	0.65	1	53	312
$\mu = 0.70$	18.90	22.00	14.03	0.74	1	45	312
CPIX Inflation							
$\mu = 0.10$	24.56	23.00	13.17	0.54	10	52	312
$\mu = 0.20$	14.29	10.00	11.76	0.82	1	41	312
$\mu = 0.30$	9.53	6.00	10.35	1.09	1	34	312
Policy Horizons Conditional on $\pi_t \notin [1.0, 3.0]$							
	Mean	Median	S.D.	C.V.	p5	p95	No. PH
CPI Inflation							
$\mu = 0.30$	40.11	46.00	18.45	0.46	9	65	72
$\mu = 0.50$	30.04	34.00	15.66	0.52	6	53	72
$\mu = 0.70$	23.90	26.00	13.52	0.57	6	46	72
CPIX Inflation							
$\mu = 0.10$	26.50	21.50	15.25	0.58	11	52	72
$\mu = 0.20$	17.62	11.50	13.24	0.75	2	42	72
$\mu = 0.30$	12.12	8.00	11.72	0.97	1	36	72

Policy horizon estimates (see text for calculation details) are for the full-sample models for CPI and CPIX inflation reported in Tables 3 and 4. The parameter μ is the tolerance criteria used to define the policy horizon as a percent deviation from the unconditional mean. The abbreviations S.D. and C.V. refer to standard deviation and coefficient of variation; p5 and p95 refer to the 5 and 95 percentile. PH refers to policy horizon.

The top panel reports summary statistics for the full set of policy horizons calculated over the 312 sample end-points 1994M1–2019M12. The bottom panel reports summary statistics for the policy horizons calculated for the 72 sample end-points for which the inflation rate is outside of the 1–3 percent inflation band.

Table 6: Policy Horizons — Sub Samples

CPI Inflation Policy Horizons Conditional on $\pi_t \notin [1.0, 3.0]$							
	Mean	Median	S.D.	C.V.	p5	p95	No. PH
2001M2–2008M1							
$\mu = 0.50$	18.65	14.00	13.32	0.71	2	45	84
$\mu = 0.50$	27.07	20.00	12.96	0.48	14	49	15
2008M2–2013M5							
$\mu = 0.50$	9.77	9.50	5.07	0.52	1	20	64
$\mu = 0.50$	10.64	9.00	5.48	0.52	5	20	25
2013M6–2019M12							
$\mu = 0.50$	2.47	1.00	2.12	0.86	1	7	79
$\mu = 0.50$	4.80	5.00	1.48	0.31	3	7	5

Policy horizon estimates (see text for calculation details) are for the sub-sample models for CPI inflation reported in Table 3. The parameter μ is the tolerance criteria used to define the policy horizon as a percent deviation from the unconditional mean. The abbreviations S.D. and C.V. refer to standard deviation and coefficient of variation. PH refers to policy horizon.

For each sub-sample, the first row are summary statistics for the policy horizons calculated over the all end-points of the sample while the second row is for the policy horizons calculated for the end-points of each sub-sample for which the inflation rate is outside of the 1–3 percent inflation band.

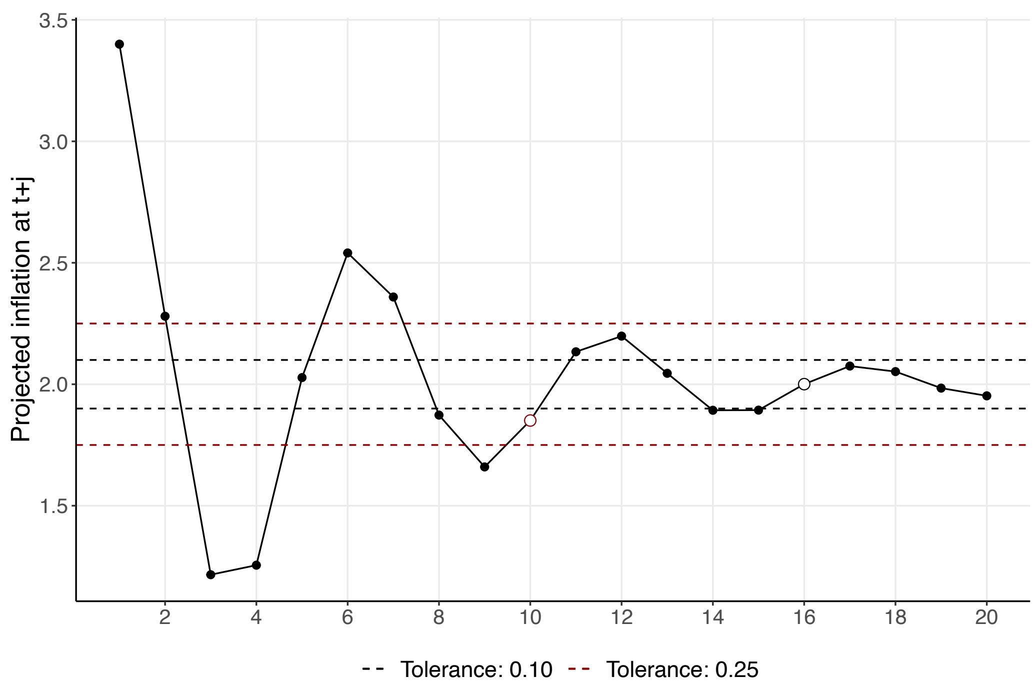

Figure 1: Conditional Inflation Forecast and Tolerance Bands

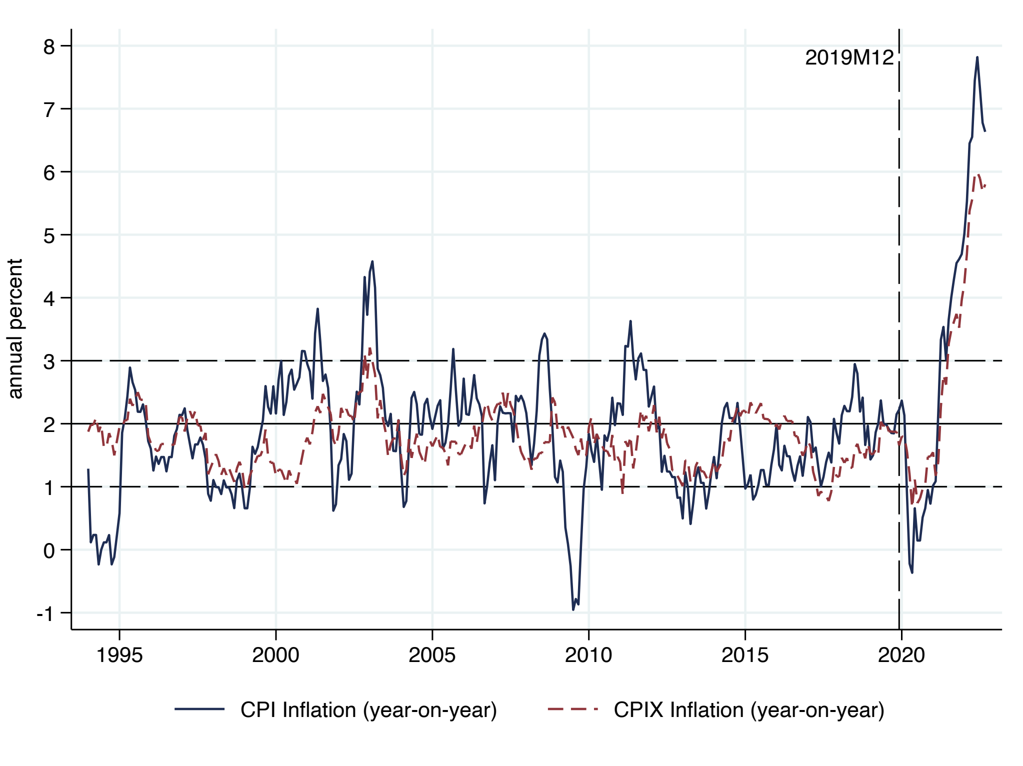

Figure 2: Inflation in Canada

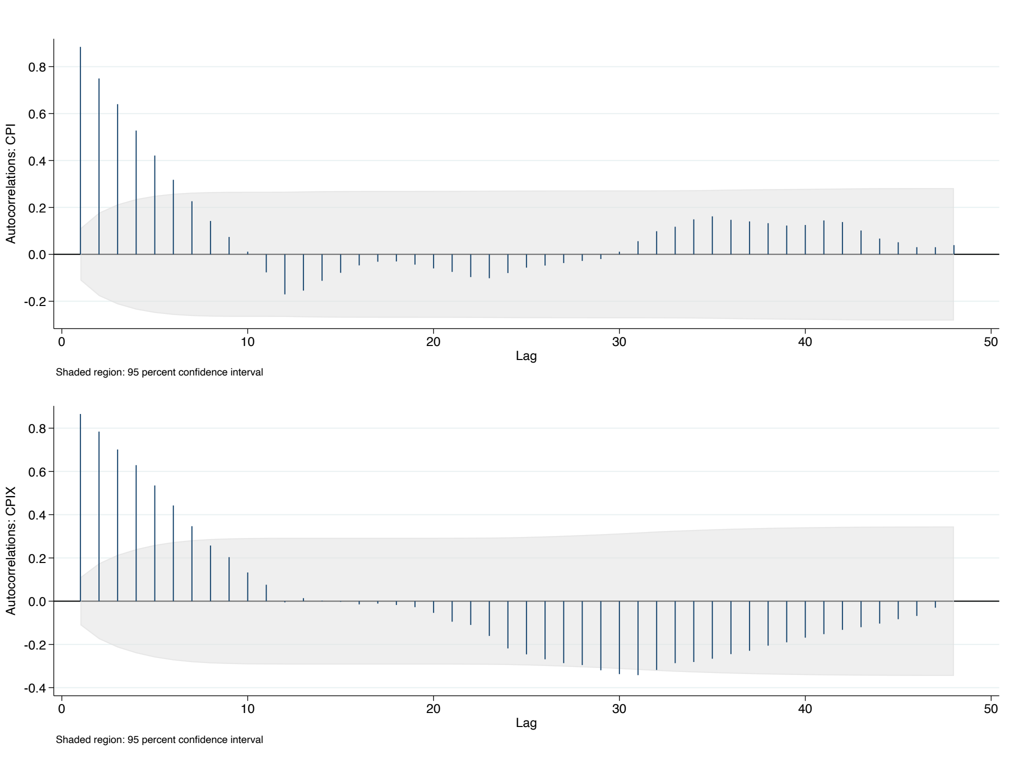

Figure 3: Inflation in Canada – Autocorrelation Functions

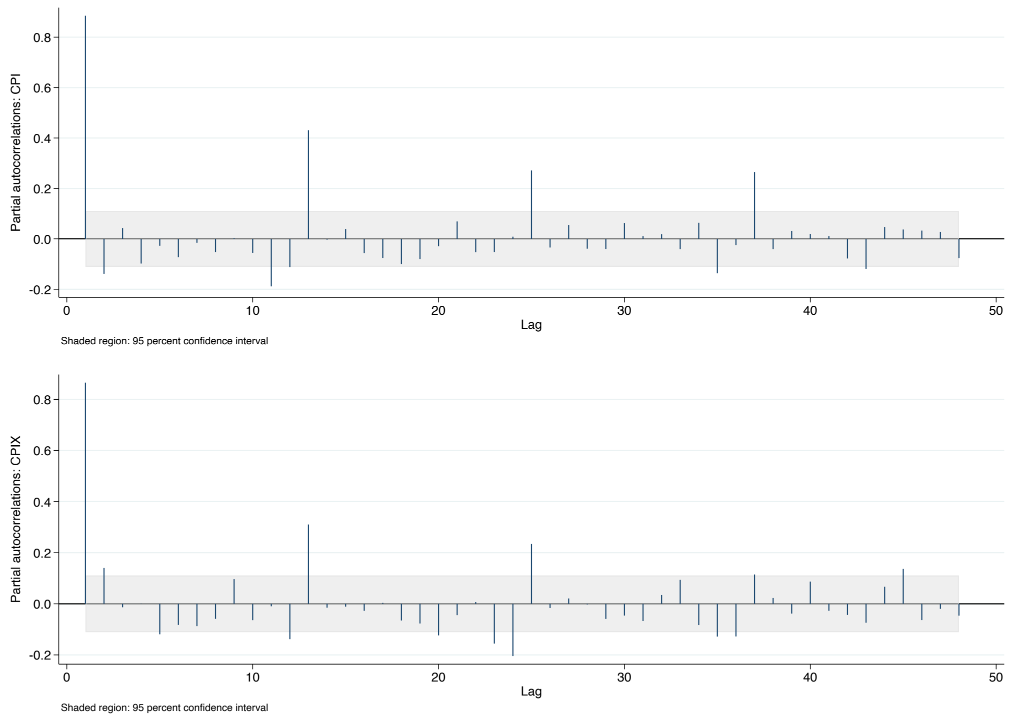

Figure 4: Inflation in Canada – Partial Autocorrelation Functions

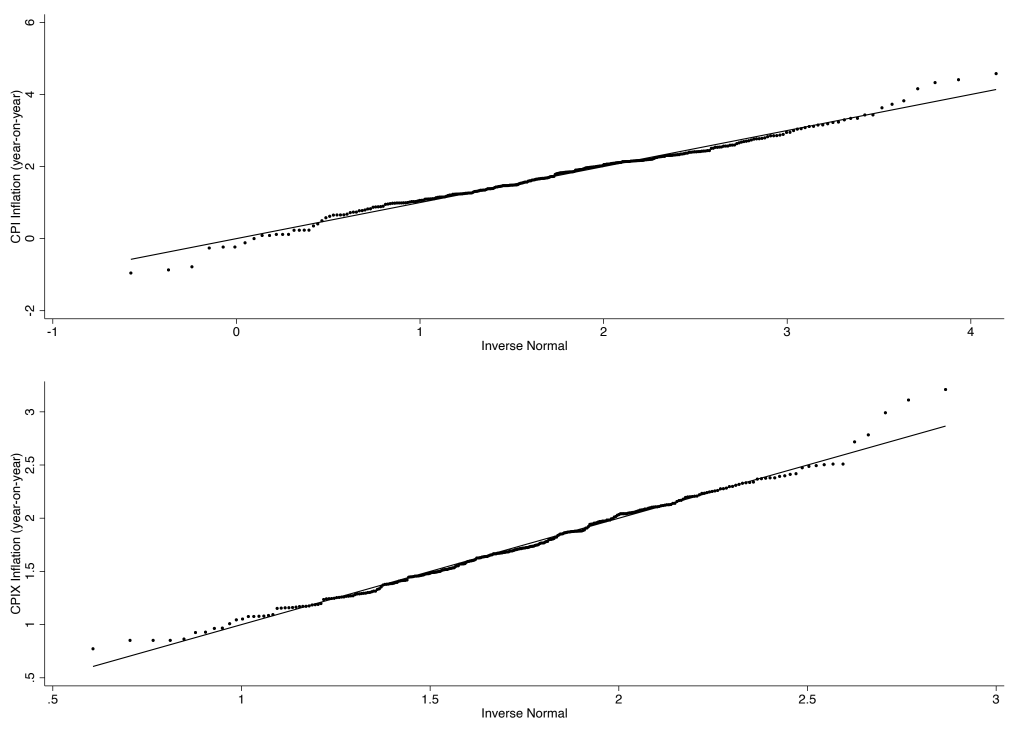

Figure 5: Inflation in Canada – QQ Plots

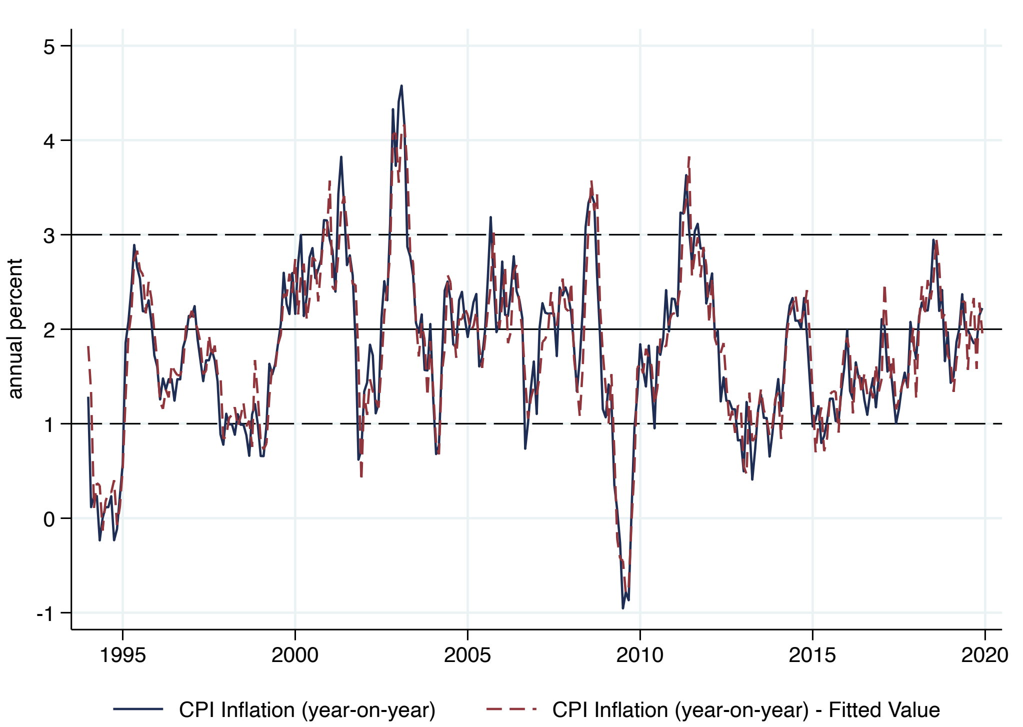

Figure 6: CPI Inflation: Actual and Fitted Values

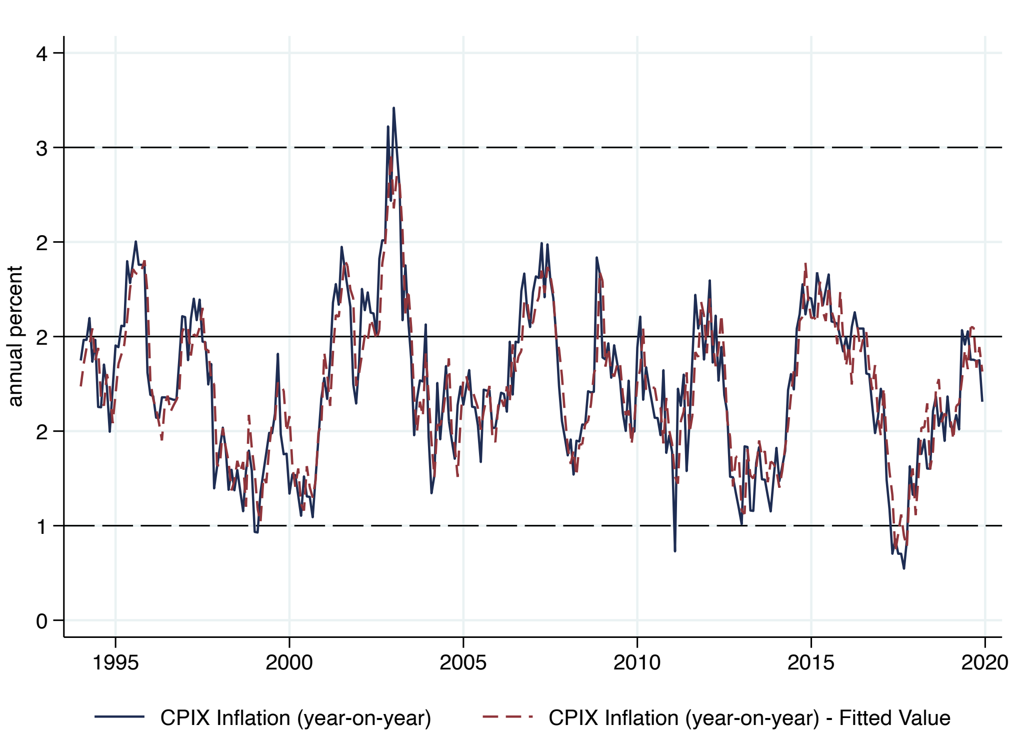

Figure 7: CPIX Inflation: Actual and Fitted Values

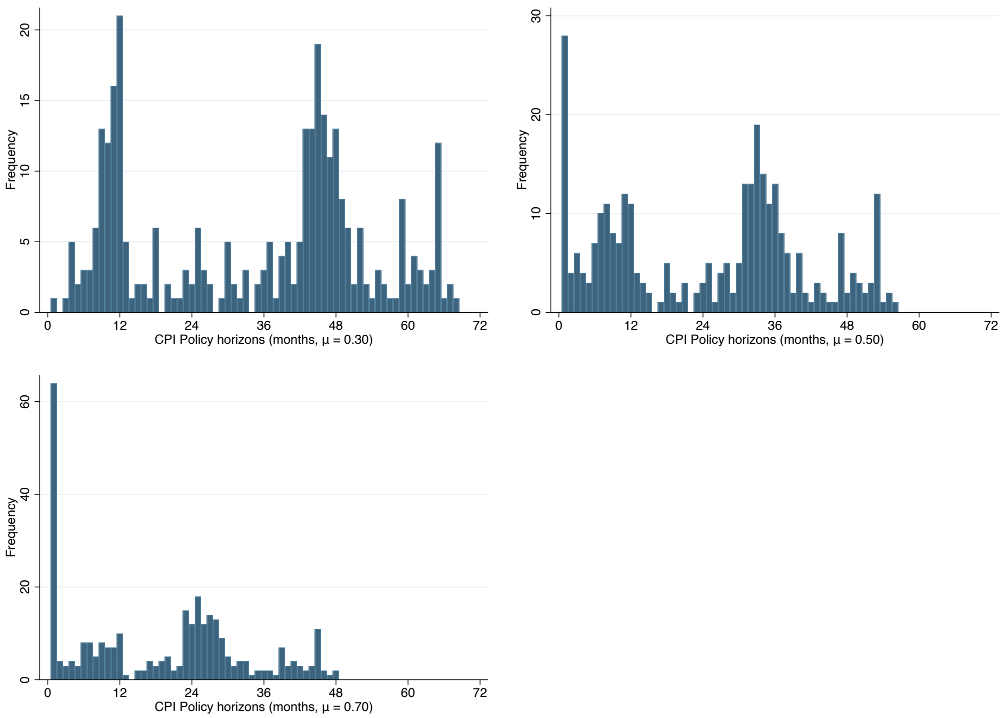

Figure 8: Full Sample Policy Horizons for CPI Inflation — All

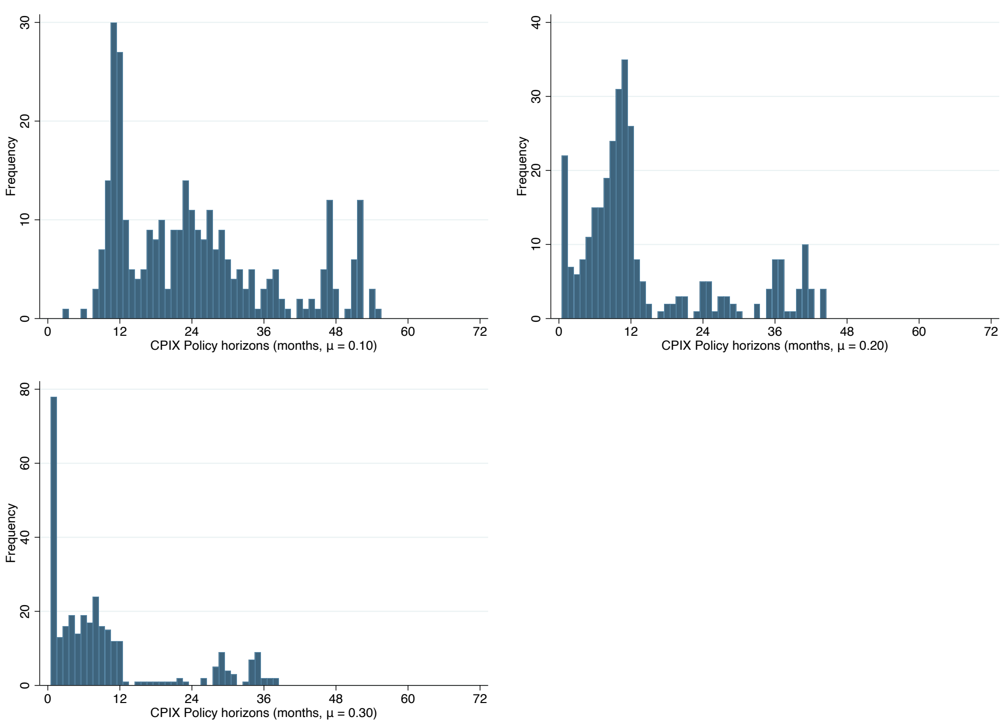

Figure 9: Full Sample Policy Horizons for CPIX Inflation - All

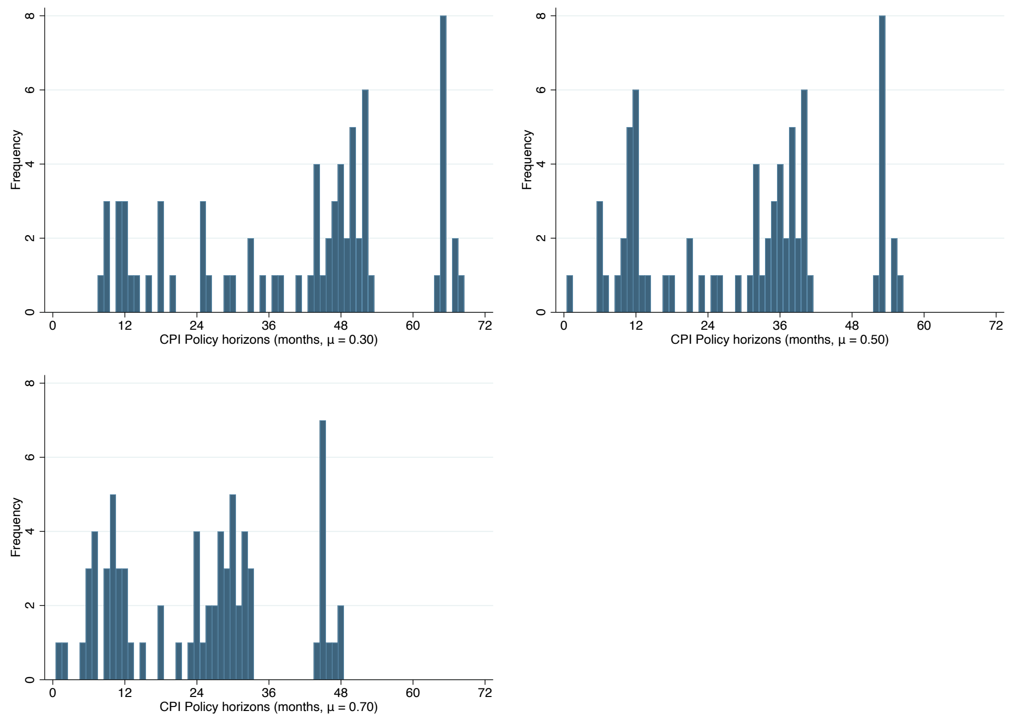

Figure 10: Full Sample Policy Horizons for CPI Inflation — $\pi_t \notin [1.0, 3.0]$

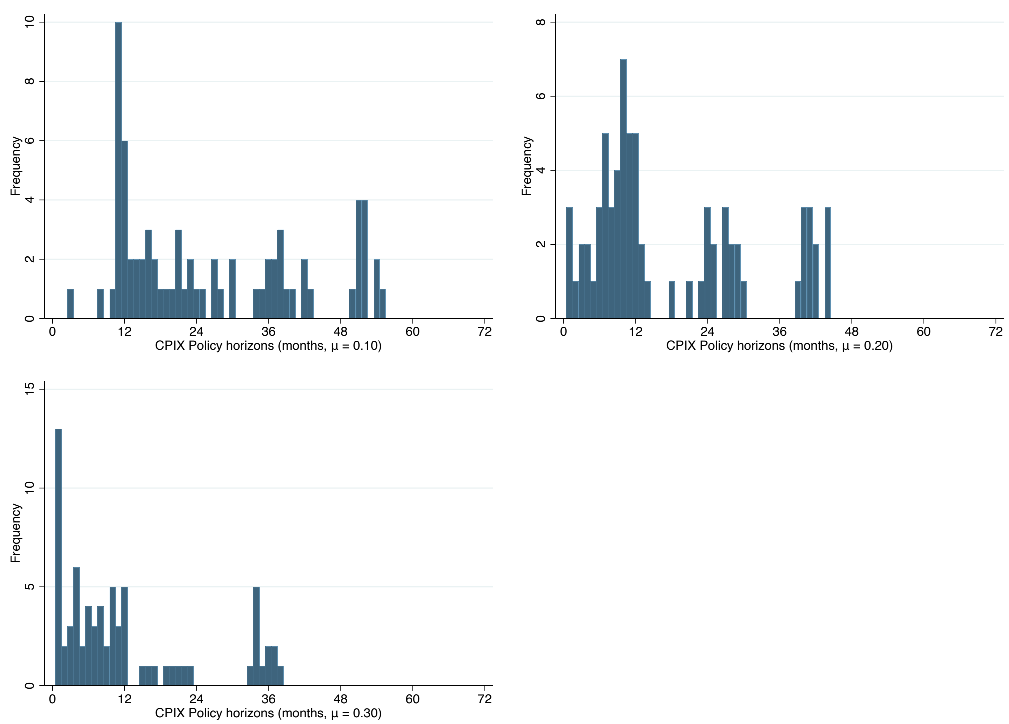

Figure 11: Full Sample Policy Horizons for CPIX Inflation - $\pi_t \notin [1.0, 3.0]$

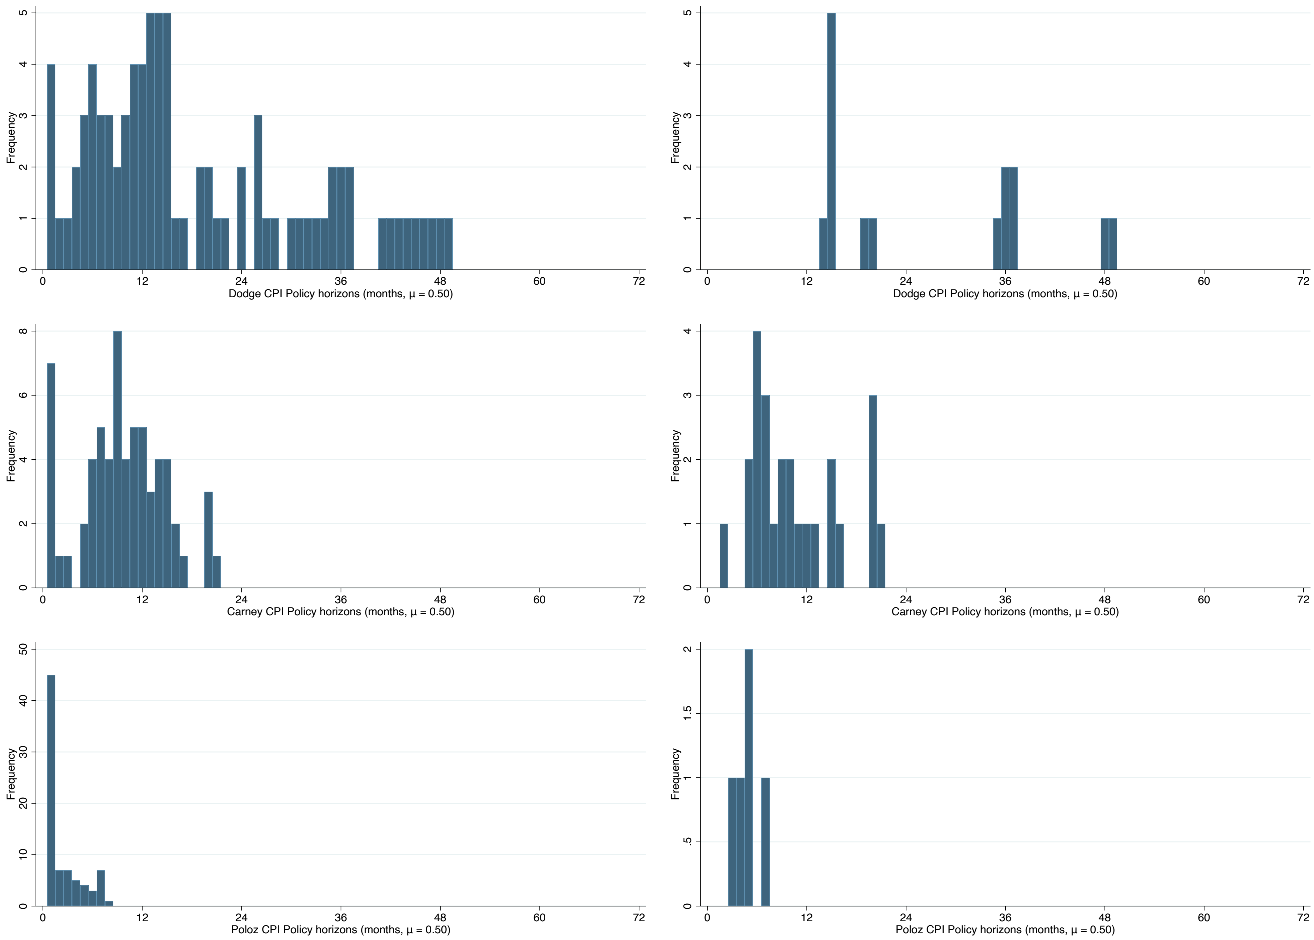

Figure 12: Sub Sample Policy Horizons for CPI Inflation - $\pi_t \notin [1.0, 3.0]$

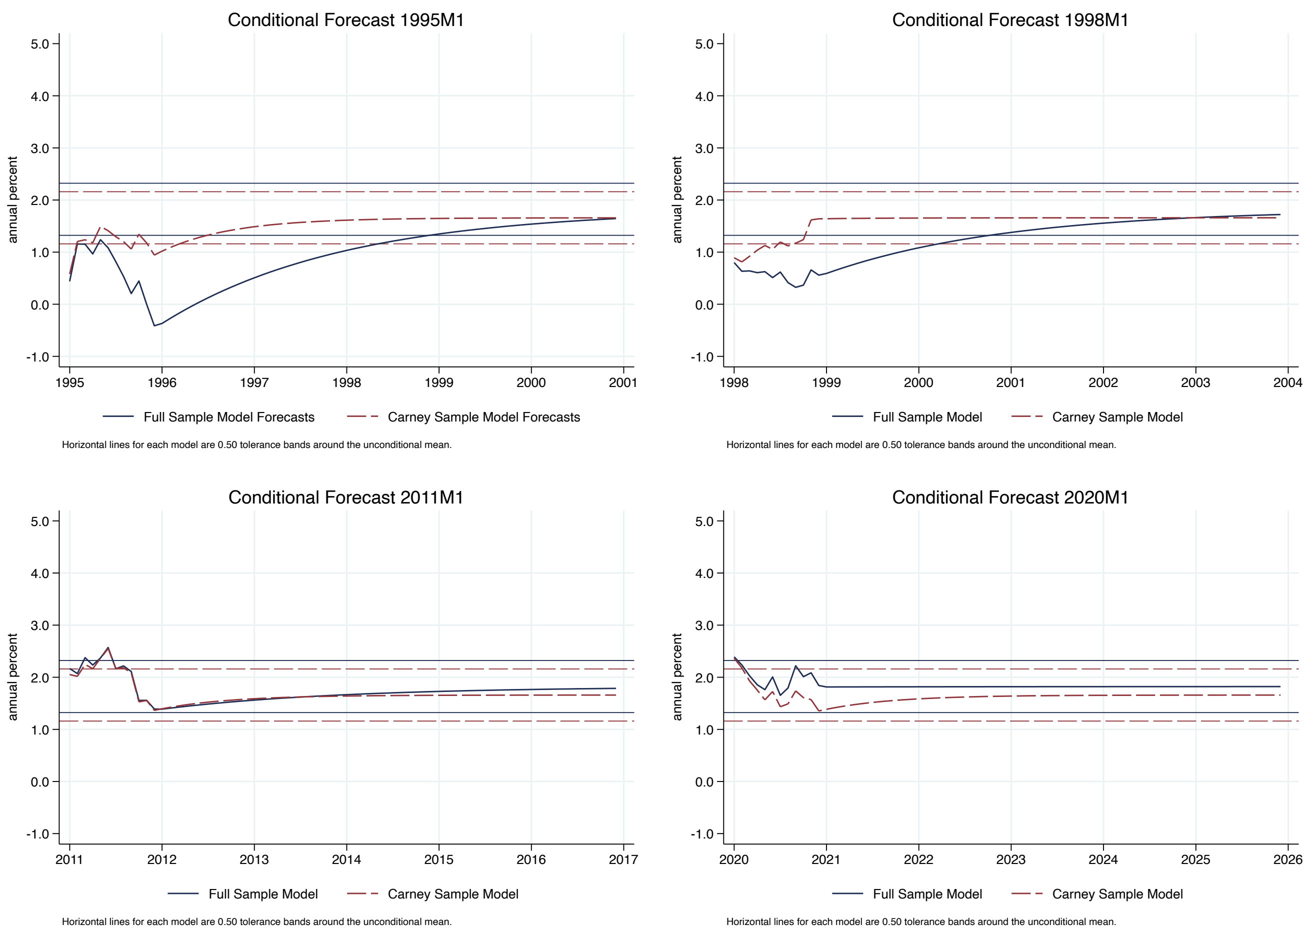

Figure 13: A Selection of Conditional Forecasts