Inflation Targeting and Policy Horizons

by

Junhong Mu
B.A., University of Victoria, 2019

A Thesis Submitted in Partial Fulfillment of the
Requirements for the Degree of

MASTER OF ARTS

in the Department of Economics

Inflation Targeting and Policy Horizons

by

Junhong Mu
B.A., University of Victoria, 2019

Supervisory Committee

Dr. Graham M. Voss, Supervisor
(Department of Economics)

Dr. Vasco J. Gabriel, Departmental Member
(Department of Economics)

Supervisory Committee

Dr. Graham M. Voss, Supervisor
(Department of Economics)

Dr. Vasco J. Gabriel, Departmental Member
(Department of Economics)

ABSTRACT

This study explores the role of flexibility in inflation targeting from two distinct perspectives. Using a qualitative approach, the first component summarizes and evaluates the monetary policy frameworks of inflation-targeting central banks across the industrialized world, focusing on the communication of monetary policy through framework documents and flagship publications. After establishing the core tenets of modern inflation targeting and key differences between the current frameworks of central banks within the OECD, I provide a detailed review of the Bank of Canada's experience with inflation targeting. I find that the role of flexibility within the Bank's policy framework has evolved in response to key events such as the 2008–09 financial crisis and the Covid-19 pandemic. The second component evaluates whether changes in the Bank of Canada's inflation targeting framework can be identified in the empirical evidence. Using conditional forecasts from a univariate model of inflation, I estimate policy horizons — the period over which forecasted inflation returns to target following a shock — from 2000 to 2022. I find that for a particular set of criteria for convergence, the range of estimated policy horizons is consistent with the Bank's stated target horizon of 6 to 8 quarters. Additionally, the distribution of estimates reflects a narrow type of flexibility in which the lengths of policy horizons vary based on the current level of inflation and the persistence of recent shocks. Estimating unique models for sample periods corresponding to the tenures of former Bank of Canada Governors Dodge, Carney, and Poloz, I find that there has been some variation in the behaviour of estimated policy horizons over time, which may reflect changes in central bank behaviour, changes in the economic environment, or a combination of both factors.

Contents

List of Tables

List of Figures

ACKNOWLEDGEMENTS

I would like to express gratitude to my partner, Bethany, who is deserving of an honourary degree in economics for the time she has dedicated to listening to rants about inflation. I am also grateful to my supervisor, Dr. Graham Voss, who wrote me countless reference letters only to get stuck dealing with me himself.

I leave Sisyphus at the foot of the mountain. One always finds one's burden again. But Sisyphus teaches the higher fidelity that negates the gods and raises rocks. He too concludes that all is well. This universe henceforth without a master seems to him neither sterile nor futile. Each atom of that stone, each mineral flake of that night-filled mountain, in itself, forms a world. The struggle itself toward the heights is enough to fill a man's heart. One must imagine Sisyphus happy.

Albert Camus

Chapter 1

Introduction

The core objective of modern central banking is the pursuit and maintenance of price stability, keeping inflation — a rise in the average price level — low and stable over time. In recent decades, inflation targeting has emerged as the prevailing monetary policy framework for central banks oriented toward such an objective.

Implicit in an inflation targeting framework is the concept of a policy horizon, the period over which monetary policy actions are expected to bring inflation back to target following a shock. Each inflation targeting central bank operates — explicitly or otherwise — under a belief of whether the appropriate policy horizon should be fixed or flexible. That is, a central bank may commit to a fixed time frame for returning inflation to target; or, it may allow the length of the horizon to vary according to how far inflation has deviated from target at a given time. This study seeks to develop a comprehensive understanding of flexible inflation targeting as currently exercised by select central banks, and subsequently estimate empirical policy horizons for Canadian inflation.

The narrative section of this study first reviews the monetary policy frameworks of inflation targeting central banks in the OECD, comparing and contrasting their respective approaches to flexibility. It then pursues an in-depth analysis of the Bank of Canada's conduct of monetary policy under inflation targeting, evaluating the changing role of flexibility in its policy framework over time.

Statements by central banks about policy horizons may not, however, be accurately reflected in the observed outcomes. The empirical section of this study proposes a simple method of estimating the *de facto* horizons for monetary policy and applies these methods in the Canadian context. Using these estimates, I assess the empirical evidence of flexibility under a much more narrow interpretation. In doing so, I also evaluate whether the estimated policy horizons are consistent with the Bank of Canada's communicated objectives. In short, I find that over the Bank's inflation targeting period, the range of policy horizons exhibits flexibility in that when inflation exceeds the 1-to-3-percent target band, larger deviations from target are generally associated with longer policy horizons. While there are distinct differences in the distribution of policy horizons over the tenures of the Bank's most recent Governors, it is not clear from the empirical evidence whether the Bank's implementation of flexibility has materially changed over this period.

The paper proceeds as follows. Chapter 2 discusses the broader issues concerning flexible inflation targets, policy horizons, and the related literature. Chapter 3 presents the sequence of narrative analysis. Chapter 4 presents the empirical analysis. Chapter 5 synthesizes findings from the preceding chapters, provides suggestions for further research, and concludes.

Chapter 2

Principles of Inflation Targeting

Pioneered by the central banks of New Zealand, Canada, and the UK, inflation targeting emerges in the early 1990s and quickly spreads to central banks around the world over the ensuing decade. While a significant portion of the initial research on this novel practice is pursued by central bankers, Lars Svensson begins to establish a theoretical basis for inflation targeting in the academic literature in 1995. In doing so, he addresses a key concern that "inflation targeting may be difficult to implement, for the simple reason that central banks have imperfect control over inflation" (Svensson, 1998). More specifically, "inflation reacts with 'long and variable lags' and with variable magnitude to changes in the monetary policy instrument. Inflation is affected by factors other than monetary policy, and sometimes with a shorter lag than monetary policy" (Svensson, 1998). This lag between a monetary policy action and the resulting effect on inflation is referred to in the early literature as the "control lag," and is conventionally regarded as a period of approximately 2 years (King, 1994; Svensson, 1999b; Woodford, 2007). In this context, monetary policy can only affect future inflation, and central banks must therefore adopt a forward-looking orientation for inflation targeting. As future inflation cannot be perfectly controlled or observed, Svensson proposes that an intermediate target be relied upon to guide monetary policy. Specifically, "the central bank's inflation forecast is indeed an ideal intermediate target: it is by definition the current variable that is most correlated with the goal, it is more controllable than the goal, and it can be made more observable than the goal" (Svensson, 1997). Therefore, the implementation of inflation targeting is more accurately described as inflation-*forecast* targeting, conditional on current information available to the central bank and an implied path for the monetary policy instrument.

The canonical New Keynesian framework establishes a formal treatment of inflation-forecast targeting, in which the central bank minimizes the discounted present value of a quadratic loss function (King, 1994; Svensson, 1999b; Svensson and Woodford, 2005):

$$E_t \sum_{\tau=t}^{\infty} \delta^{\tau-t} \mathcal{L}(\pi_\tau) \tag{2.1}$$

In its simplest form, the period loss function $\mathcal{L}(\pi_t)$ penalizes the squared deviation of inflation from the long-run inflation target, π^*. Svensson (1997) alters Eq. (2.1) to represent inflation-forecast targeting:[1]

$$E_t \sum_{\tau=t}^{\infty} \delta^{\tau-t} \mathcal{L}(\pi_{\tau+h}) \tag{2.2}$$

where $\mathcal{L}(\pi_{\tau+h})$ evaluates the deviation of forecasted inflation from the inflation target. From the first-order condition, Svensson (1997) then derives a simple criterion for inflation-forecast targeting:

$$E_t(\pi_{t+h} - \pi^*) = 0 \tag{2.3}$$

That is, the central bank sets its policy rate at time t such that the conditional h-period-ahead forecast for inflation achieves the inflation target.

Eq. (2.3) reflects a "strict" inflation target, under which the central bank's sole consideration is returning inflation to target as quickly as possible. However, strict inflation targeting may at times require an aggressive monetary policy response that exacerbates output and real exchange rate volatility (Svensson, 1997). Woodford (2004) advocates for a target criterion that explicitly considers real variables in addition to inflation, allowing "some degree of temporary variation in the inflation rate in response to real disturbances, for the sake of greater achievement of other stabilization objectives." Specifically, the target becomes "flexible" in that it is weighted between stabilizing inflation in the long run and minimizing the short-run volatility in the real economy (Woodford, 2004). For the simple case where the only additional variable is the deviation of the output gap from a targeted value (usually zero), the loss function becomes $\mathcal{L}(\pi_\tau, x_\tau)$. The associated first-order condition is:

$$E_t(\pi_{t+h} + \lambda x_{t+h} - \pi^*) = 0 \tag{2.4}$$

where x_t is the output gap at time t and λ is the associated weight relative to inflation. In general, the complementary objective x_t may be replaced or augmented with the real exchange rate, full employment, interest rate volatility, or any other variable that may be pertinent to the central bank's mandate (Svensson, 1999a; Goodfriend, 2003; Coletti et al., 2006).

Within this structural framework, two broad classes of policy rules are consistent with in-

[1] For the remainder of this chapter, any reference to inflation targeting will pertain specifically to inflation-*forecast* targeting.

flation targeting (Svensson, 2002). The first is a simple *instrument rule*, an explicit reaction function that maps the monetary policy instrument — usually a nominal interest rate — to a limited set of predetermined variables. A well-known example is a Taylor-type rule in which the policy rate i_t is determined as a linear combination of current inflation π_t, its deviation from the inflation target π^*, the output gap x_t, and the 1–period lagged interest rate i_{t-1} (Svensson, 2002). This type of rule allows the instrument path to be easily anticipated by the private sector, but is likely not consistent with the actual practice of inflation targeting — no central bank would reasonably commit to a rule in which its policy decisions are mechanically determined by a small set of variables. For this reason, instrument rules such as the Taylor rule are better viewed as a "guide" for future decisions from which the policy rate may deviate as needed in response to pertinent economic conditions (Taylor, 1993).

In comparison, a *targeting rule* assigns target variables — more accurately, their forecasted values — and their target levels over a loss function such as Eq. (2.4) (Svensson and Woodford, 2005). As the loss function implies an optimal interest rate in each period, the optimal sequence of interest rates can be viewed as an implicit instrument rule or an endogenous reaction function of both endogenous and exogenous variables (Rudebusch and Svensson, 1999). As adherence to a targeting rule is contingent on the observed outcome of the target variable rather than the path of the policy instrument, the functional form of this implicit rule will "generally change in the case of changes in the model of the economy used in implementing the rule" (Svensson and Woodford, 2005). For this reason, a targeting rule offers greater flexibility for monetary policy to accommodate real-world shocks than an instrument rule, and in the view of Svensson (1999a) and Rudebusch and Svensson (1999), more accurately depicts a central bank's decision-making framework under inflation targeting.

In their assessment of optimal policy rules, Svensson and Woodford (2005) find that targeting rules offer more transparency than explicit instrument rules. Evaluated on the basis that "the connection between the central bank's decision process and its ultimate objective [should be] as transparent as possible," targeting rules are specified in terms of the desired behaviour of (forecasted) inflation relative to its target; in comparison, explicit instrument rules are specified in terms of current or lagged variables and disturbances which are less clearly related to the inflation targeting objective (Svensson and Woodford, 2005). This assessment emphasizes that transparency should concern the policy objective rather than the policy instrument, a view that is supported by Woodford (2007) — "the targeting approach should make the consequences of the central bank's decisions more predictable, since one should be able to count on the target variables satisfying the target criterion to a reasonable extent, even if the path of the policy rate that this involves will not always be highly predictable." In practice, however, central banks rarely communicate in terms of policy rules. Instead, they rely on serial publications — for example, the Bank of Canada's quarterly *Monetary Policy Report* — to provide information about the policy-making process and establish expectations for the near future (Woodford, 2004). Transparency, therefore, relies upon observable benchmarks by which the public can evaluate the central bank's performance relative to its stated goal. In this context, Davig and Foerster (2021) uses a simple structural framework to illustrate that by specifying a point target for inflation, a tolerance band around the target, and an inflation forecast, the central bank can communicate the same information as if it revealed a specific policy rule. By providing an actionable objective for the policy rule — for example, that "inflation will be within the tolerance band in N periods" — the policy horizon N serves as a measure of accountability to the public (Davig and Foerster, 2021).

The policy horizon — the number of periods over which inflation is expected to return to target — shares conceptual similarities with the forecast horizon h from Eqs. (2.3) and (2.4), as both pose the question of how forward-looking monetary policy should be. With a strict inflation target, Svensson (1997) considers "a horizon corresponding to the control lag" to be optimal as it is the shortest possible horizon over which monetary policy affects future inflation. This implies that, in Svensson's original framework, the forecast horizon and the policy horizon are one and the same. Using data from the US and the UK for 1953 to 2001, Batini and Nelson (2001a) estimate the control lag to be between 1 and 3 years. This is consistent with Svensson (1997), who proposes lags of 1.5 years for strict inflation targeting and 2.5 years for flexible inflation targeting.

Batini and Nelson (2001b), however, argues that policy horizons are conceptually distinct from the control lag of monetary policy. To differentiate, the authors offer two distinct definitions of the optimal horizon for inflation targeting, corresponding to the two classes of policy rules discussed above. First, consider an explicit instrument rule for the nominal interest rate:

$$i_t = c + \beta[E_t\pi_{t+h} - \pi^*] \tag{2.5}$$

where $E_t\pi_{t+h}$ is the conditional forecast at time t of inflation h periods ahead, π^* is the inflation target, β is the central bank's discount factor, and c is a constant. In this context, the horizon h reflects the central bank's view of how far ahead it should forecast inflation to "bypass the transmission lags" of monetary policy (Batini and Nelson, 2001b). Given a specific model of the economy, there is a unique optimal horizon h at which to forecast inflation to minimize the loss function subject to the instrument rule Eq. (2.5). In this way, the horizon is a policy choice and an optimized parameter, and therefore lacks flexibility in that it does not allow for a range of optimal values. Batini and Nelson (2001b) denote this the optimal *feedback horizon* (OFH).

In comparison, given a targeting rule for inflation, Batini and Nelson (2001b) define the optimal *policy horizon* (OPH) as "the number of periods it takes for inflation to settle on target after a shock, under the optimal rule for the instrument."[2] Rather than an optimized parameter as in the case of the previous definition, the optimal policy horizon is "a metric associated with an optimal rule," and may therefore take on a range of values (Batini and Nelson, 2001b). Using impulse response functions from VAR and semi-structural models, Batini and Nelson (2001b) shows that the length of the optimal policy horizon varies depending on the type of shock hitting the economy. This variation reveals a type of flexibility that is much closer to the view of flexibility taken by modern central banks than Svensson's interpretation of a weighted target between inflation and the output gap as in Eq. (2.4). For example, in its July 2022 *Monetary Policy Report*, the Bank of Canada states that "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, 2022b). This approach implies that while the central bank is sensitive to the impact of its policy decisions on the real economy, it does not pursue an explicit target for any particular measure. Instead, trade-offs between the inflation target and certain components of the real economy are context-dependent and may necessitate flexibility in the lengths of policy horizons over time. As precisely summarized by Coletti et al. (2006):

> "A short horizon would be consistent with a vigorous change in interest rates in order to return inflation to target quickly, but could result in excessive volatility in interest rates and the real economy, since the lagged effects of vigorous interest rate changes need to be cancelled by subsequent actions in the other direction. A long horizon would be consistent with a more sluggish change in interest rates that could result in less real volatility, but would cause deviations of inflation from target to be more persistent. Thus, there is an optimal inflation-target horizon that balances these two opposing considerations."

To formalize this interpretation of flexibility, consider the following modification to Eq. (2.4):

$$E_t(\pi_{t+h} - \pi^*) = 0 \qquad \forall h \geq \bar{h}(E_t s_{t+h}) \tag{2.6}$$

where $\bar{h}_t$ is the optimal policy horizon and s_t is a state variable. Note that the time subscript on $\bar{h}_t$ implies that the optimal horizon need not be constant over time. The state variable s_t captures an arbitrarily large set of variables that may determine inflation — and its deviation from target — at any given time. For example, Batini and Nelson (2001b) evaluates the impact of shocks from three sources — aggregate supply, aggregate

[2] As Svensson (1999b) and Rudebusch and Svensson (1999) illustrate, an explicit instrument rule can be solved from the first-order condition of a central bank's loss function. Furthermore, Svensson and Woodford (2005) shows that under particular specifications, the instrument rule implied by a targeting rule is consistent with "the optimal stationary equilibrium" of the canonical New Keynesian model.

demand, and the exchange rate — all of which can be viewed as a particular realization of a state of the economy. In practice, however, inflation may be determined by factors including but not limited to the recent history of inflation, domestic indicators such as housing prices, household balance sheets, unemployment, proximity to the effective lower bound for monetary policy, and the behaviour of the central bank itself, as well as their respective interactions. While it is neither feasible nor intuitive to specify each of these variables in an explicit policy rule or a weighted inflation target, a state variable allows any combination of variables to be captured at a given time. In addition, since the inflation target is forward-looking, $\bar{h}$ is determined not by the current state of the economy, but by the central bank's expectation of how the state of the economy will evolve over its forecast horizon. This is denoted in Eq. (2.6) as $E_t s_{t+h}$.

Eq. (2.6) generalizes the concept of flexibility exhibited in the range of policy horizons estimated by Batini and Nelson (2001b), which at its core need not be founded upon a targeting rule or imply any notion of optimality. Borrowing from the language of central banks' policy frameworks, it is simpler to consider $\bar{h}$ as the "appropriate" or "desired" horizon over which inflation is expected to return to target. In this context, Eq. (2.6) may be interpreted as "conditional on the expected path of the economy, the central bank sets the policy rate such that its forecast for inflation achieves the inflation target within $\bar{h}$ periods."

This interpretation is advantageous in that it integrates the two most relevant applications of flexibility in the literature — first, that a central bank may be flexible in that its commitment to an inflation target is weighed against stabilizing other variables it deems relevant; second, that it may be flexible in that the appropriate horizon to return inflation to target is allowed to vary over time. This view of flexibility is much more general than Svensson's flexible inflation target as in Eq. (2.4). For this reason, I denote the former as *broad* flexible inflation targeting and the latter as *narrow* flexible inflation targeting. In both the narrative review in Chapter 3 and the empirical analysis in Chapter 4, I will rely on the broad definition. However, given the limitations of my empirical framework, only a specific application of the broad definition will be relevant to the associated analysis. I define this application in detail in Section 4.1.

Chapter 3

Narrative Analysis

3.1 Flexibility in OECD Central Banks

3.1.1 Inflation Targets and Policy Horizons

The narrative analysis seeks to establish a broad understanding of the inflation targeting frameworks utilized in modern industrialized economies, for which the member countries of the Organisation for Economic Co-operation and Development (OECD) are natural candidates. In an effort to review the central banks of all 38 OECD members — with 17 collectively represented by the European Central Bank — four institutions have been omitted. I exclude Colombia's Banco de la República and the Banco Central de Costa Rica due to language restrictions on the pertinent policy documents. I exclude the Danmarks Nationalbank as its monetary policy framework focuses on exchange rate management rather than inflation targeting, and therefore falls outside the scope of this study. I also exclude the Central Bank of the Republic of Turkey due to its ongoing experience with significant political interference, which has undermined the validity of its official policy framework.

The remaining 18 central banks are listed and categorized in Table 3.1 below. In the style of Table 1 in Davig and Foerster (2021), I sort the current inflation target of each central bank into four distinct approaches based on the specification of a point target versus a target range and the inclusion of a tolerance band. As shown in Table 3.1, 15 of 18 central banks identify a point target for inflation (columns (a) and (b)), but only seven of 15 then specify a tolerance band around the target. Of the three remaining central banks (columns (c) and (d)), two specify a target range without a midpoint while the other aims to keep inflation below an upper bound.

Column (a) of Table 3.1 reflects *uncertainty ranges* that "inform the public about the [central bank's] assessment of the magnitude of inflation variations under appropriate policy"

Table 3.1: List of Inflation Targets

(a) Band w/ Mid		(b) No Band		(c) Band w/o Mid		(d) Upper Bound	
Canada	2 ±1	Euro Zone	2	Australia	2–3	Switzerland	< 2
Chile	3 ±1	Iceland	2.5	Israel	1–3		
Czech Republic	2 ±1	Japan	2				
Hungary	3 ±1	Norway	2				
Mexico	3 ±1	South Korea	2				
New Zealand	2 ±1	Sweden	2				
Poland	2.5 ±1	UK	2				
		USA	2				

Note: List of inflation targets compiled for 2022 by Central Bank News.

(Chung et al., 2020). Under this interpretation, tolerance bands provide monetary policy with additional flexibility, as "due to shocks and uncertainties, no central bank is expected to hold inflation exactly at the target month after month, year after year" (Andersson and Jonung, 2017). Given sufficient credibility, central banks do not need to view small fluctuations within the band as breaches of the target. Under inflation targeting frameworks that consider additional objectives such as financial stability and real activity, this flexibility allows central banks to pursue alternative policy paths that may bring inflation back to target over a longer horizon.

The Bank of Canada's implementation of a 1-to-3–percent tolerance band — which it calls an inflation-control range — around its 2–percent target serves as "a reflection of the short-run uncertainty of outcomes stemming from unpredictable shocks and not as a measure of the indifference of the Bank as to the outcome" (Bank of Canada, 2001). This range allows the Bank to exercise a degree of flexibility, "looking through" minor fluctuations in inflation to mitigate the potential impacts of interest rate changes on the real economy (Bank of Canada, 2016b). Specifically, while the Bank typically aims to return inflation to 2 percent within 6 to 8 quarters, "the most appropriate horizon for returning inflation to target will vary depending on the nature and persistence of the shocks" (Bank of Canada, 2022c). In the latest update to its policy framework in December 2021, the Bank also emphasizes that it will leverage the flexibility of the inflation-control range to "actively seek the maximum sustainable level of employment, when conditions warrant" (Bank of Canada, 2021e). I provide an in-depth discussion of the details, execution, and evolution of the Bank of Canada's inflation targeting framework in Section 3.2.

In comparison, the Reserve Bank of New Zealand's (RBNZ) monetary policy framework allows for a much greater degree of flexibility. Alongside inflation, the Bank considers employment and output as indicators of the real economy, conducts policy to support maximum sustainable employment, and, as of its *Monetary Policy Committee Remit* from 2021, also assesses the effects of its monetary policy on sustainable housing prices (Reserve Bank of New Zealand, 2021). In this context, the inflation target is initially established as a target range of 1 to 3 percent over the medium term, giving the Bank sufficient manoeuvrability to address this wide array of policy considerations while also accommodating short-term deviations from target (Reserve Bank of New Zealand, 2022). A 2–percent midpoint is added in 2012 as a means to anchor inflation expectations, as "without a mid-point target, inflation expectations may float to either end of the 1–percent to 3–percent band, depending on the economic outlook" (Reserve Bank of New Zealand, 2022). Similar to the Bank of Canada, the RNBZ sets policy "with the aim of returning inflation to the target mid-point in about 2 years, and allowing this time horizon to vary depending on the shocks hitting the economy" (Reserve Bank of New Zealand, 2020).

The remaining central banks in column (a) generally operate in manners similar to the Bank of Canada. The Banco Central de Chile, however, stands out as the only central bank that pursues a fixed policy horizon. Specifically, "the Bank conducts its monetary policy seeking that, irrespective of the current level of inflation, its forecast for a 2–year horizon will be 3 percent" (Banco Central de Chile, 2022). In other words, while a ± 1 percentage point tolerance band accommodates short-run fluctuations in inflation, the Bank must pursue a policy path that fully stabilizes forecasted inflation at the 3–percent target within 2 years. In doing so, the Bank forgoes the flexibility associated with longer policy horizons, choosing instead to reduce uncertainty and better anchor inflation expectations around its point target.

In the absence of a tolerance band, the central banks in column (b) of Table 3.1 generally place more emphasis on the medium-term orientation of their inflation targets to accommodate short-term volatility. In doing so, a trade-off arises between an elevated level of uncertainty around the inflation target in the short term and the greater degree of flexibility granted by a less restrictive target. Due to this flexibility, there is notably more variation in the monetary policy frameworks of these central banks compared to those with a tolerance band.

For the European Central Bank (ECB), this orientation "allows for inevitable short-term deviations of inflation from the target, as well as lags and uncertainty in the transmission of monetary policy to the economy and to inflation" (Lane, 2022). As it considers economic growth, full employment, financial stability, and inflation as "mutually consistent objectives," the ECB necessarily relies upon flexible policy horizons as a core feature of its medium-term orientation (European Central Bank, 2021). This is established in the ECB's monetary policy strategy, which states that "the Governing Council may decide to lengthen the horizon over which inflation returns to the target level in order to avoid pronounced falls in economic activity and employment" (European Central Bank, 2021). Although they do not each share the same set of complementary objectives enumerated by the ECB, the central banks of Norway, South Korea, and Japan operate under very similar approaches to medium-term inflation targeting and flexible policy horizons.

Contrary to the RBNZ, which added a midpoint to its original target range for inflation, Sweden's Sveriges Riksbank elected to remove the ± 1 percentage point tolerance band from its 2–percent inflation target. The tolerance band was initially viewed as a "pedagogical tool intended to illustrate that certain deviations [...], while not being too large, had to be accepted" (Sveriges Riksbank, 2016). By 2010, however, the Riksbank considered it obsolete, as "there is considerable understanding for the fact that inflation commonly deviates from the target and that the deviations are sometimes larger than 1 percentage point. Inflation can thus be outside of the tolerance interval without threatening the credibility of the inflation target. Such deviations have proved to be a natural part of monetary policy" (Andersson and Jonung, 2017). However, in recent years, the Riksbank has revisited the merits of that assessment and signalled toward the reintegration of some sort of uncertainty band in the future. Despite the differences in their policy frameworks, the Riksbank's approach to flexible policy horizons is similar to that of the Bank of Canada and the RBNZ. Specifically, while "the Riksbank's general ambition has been to adjust monetary policy so that inflation is expected to be fairly close to the target in 2 years' time," it recognizes that "there is no general answer to the question of how quickly the Riksbank aims to bring the inflation rate back to 2 percent if it deviates from the target" (Sveriges Riksbank, 2021).

While the Bank of England and the Central Bank of Iceland do not specify tolerance bands in the sense of an uncertainty range, they offer a different application that relates to the transparency of their communications. To reinforce inflation expectations around their respective 2–percent targets when inflation experiences significant short-run volatility, both central banks are required by legislation to submit an open letter should inflation breach a buffer zone around the target. Should inflation deviate from target by more than 1 percentage point, the Bank of England must address the Chancellor with an explanation of why inflation has deviated from target, the Bank's policy response, and "the horizon over which the Committee judges it is appropriate to return inflation to the target" (Bank of England, 2013). The Central Bank of Iceland is granted a slightly larger buffer of ± 1.5 percentage points, but it is similarly required to submit a letter to the Government explaining the reasons for deviation, outlining the planned course of action, and establishing the horizon over which it intends for inflation to return to target (Central Bank of Iceland, 2001). In both cases, the policy horizons are implied to be flexible as the determination of the optimal length is delegated to the central bank without explicit restrictions.

The US Federal Reserve stands out as an outlier among the eight central banks that do not specify a tolerance band. Since its framework review in 2020, the Fed has departed from its flexible inflation target — similar to that of the ECB — in favour of *average* inflation targeting. In its monetary policy strategy, the Federal Open Market Committee (FOMC) states that "in order to anchor longer-term inflation expectations at [2 percent], the Committee seeks to achieve inflation that averages 2 percent over time, and therefore judges that following periods when inflation has been running persistently below 2 percent, appropriate monetary policy will likely aim to achieve inflation moderately above 2 percent for some time" (Board of Governors of the Federal Reserve System, 2020). The core motivation for this change arises from the Fed's recent experience with low inflation, and the potential for downward pressure on inflation when Fed funds rate is already at its effective lower bound. Through this policy shift, the Fed believes that it will be able to more effectively anchor inflation expectations and manage its dual mandate between price stability and maximum employment. In particular, "the concern is that below-target inflation outcomes may become entrenched in lower, below-target inflation expectations. [...] Average inflation targeting allows policy space for the Fed to 'make up' for lost inflation to sustain the [...] equilibrium" (Martínez-García et al., 2021). This capacity to "make up" inflation marks the crucial difference between flexible and average inflation targeting. In the former case, the inflation target is forward-looking in that the central bank does not take past deviations from target into account when determining the optimal policy path — letting "bygones be bygones". In contrast, average inflation targeting is backward-looking, allowing inflation to "modestly and temporarily run above the target to make up for past shortfalls, or vice versa" (Martínez-García et al., 2021). Nevertheless, this policy shift does not alter the use of flexible policy horizons, which allows the Fed to consider the appropriate horizon for returning both inflation and employment to their respective targets when the two objectives are not complementary (Board of Governors of the Federal Reserve System, 2020).

Each of the three central banks in columns (c) and (d) of Table 3.1 operate without a point target for inflation altogether. The Reserve Bank of Australia (RBA) seeks to hold inflation between 2 to 3 percent "on average, over time" (Reserve Bank of Australia, 2016). Despite the syntactic similarities to the Federal Reserve, the RBA clarifies that its inflation target is necessarily forward-looking, allowing "a role for monetary policy in dampening the fluctuations in output over the course of the cycle" while also providing the flexibility to consider financial stability (Reserve Bank of Australia, 2022). Referring to this target range as a "thick point," Chung et al. (2020) clarifies that it does not represent "a zone of policy inaction," but rather reflects the fact that the measurement of inflation and inflation forecasts cannot be perfectly precise and there is little to be gained by pursuing that precision. Although the RBA's monetary policy framework omits any mention of the appropriate policy horizon, its emphasis on a medium-term orientation and the flexibility around its inflation target indicate that at minimum, it is not compelled to return inflation to target within a fixed time frame.

In comparison, the Bank of Israel sets a slightly wider target range of 1 to 3 percent but does not specify a medium-term orientation for monetary policy. Also absent are explicit considerations for policy objectives other than price stability. In this context, the Bank "sets the rate of interest at a level that will keep inflation within the target range, or that will return inflation to within the range in a period not longer than 2 years" (Bank of Israel, 2014). Although their policy frameworks establish different approaches to the inflation target, the Bank of Israel and the Banco Centrale de Chile are the only two central banks in the sample that specify a fixed policy horizon. In both cases, the inflation target itself can be interpreted as "strict" in the sense discussed in Section 2 — that is, considerations such as the real economy and financial stability do not fall within the scope of monetary policy, whose sole focus is inflation.

Finally, the Swiss National Bank (SNB) is the only central bank in the OECD that targets an upper bound for inflation, equating "price stability with a rise in the Swiss CPI of less than 2 percent per annum" (Swiss National Bank, 2022a). Given that "deflation, i.e. a sustained decrease in the price level, also breaches the objective of price stability," the upper bound could be reasonably interpreted as a target range of 0 to 2 percent in practice (Swiss National Bank, 2022a). However, this lower bound is not explicit and the SNB's medium-term view of inflation targeting generally allows monetary policy to ignore one-off factors that may push inflation above 2 percent or below 0 percent. This approach provides the SNB with a significant degree of flexibility, taking into consideration "a large range of indicators of domestic and international economic and monetary developments, as well as of financial stability, for its monetary policy decisions" (Swiss National Bank, 2022a). While the SNB relies upon conditional forecasts over a 3–year horizon to ensure that its monetary policy is "appropriate for maintaining price stability," its mandate does not require that inflation be returned to within the target range over that horizon. Specifically, "the SNB does not react mechanically to the conditional inflation forecast," implying that, at least within a 3–year horizon, the policy horizon may vary in response to one-off shocks (Swiss National Bank, 2022a).

The wide range of approaches explored in this section reveals several common themes. First, given the transmission lag of monetary policy and the natural uncertainty associated with inflation, there is a common emphasis on the medium-term orientation of inflation targeting. While few central banks reference the term explicitly, this emphasis is consistent with the definition of inflation-*forecast* targeting from Section 2. Second, the majority of central banks are not strict inflation targeters, acknowledging to some extent the role of monetary policy in objectives such as financial stability, output, and employment. Third, despite differences in economic conditions and policy frameworks, all but two of the central banks in the sample operate with an understanding that the appropriate policy horizon should be flexible. There is, however, significant variation in the extent and specificity to which those policy horizons are communicated in practice. The following section explores central banks' communication strategies in detail.

3.1.2 Communication and Transparency

Most central banks identify the appropriate policy horizon through conditional forecasts for inflation, from which one can deduce the number of periods over which inflation is expected to converge to target. Generally, the forecasts are reported and updated in flagship publications such as the Bank of Canada's quarterly *Monetary Policy Report* (or equivalent), in which the central bank assesses the prevailing economic conditions and key trends, and communicates its outlook for the economy.

The amount and quality of information that can be extracted from inflation forecasts can be viewed as a measure of transparency for the central bank's conduct of monetary policy. Each central bank's communication of its inflation outlook may be disaggregated into four general components — (1) convergence toward the inflation target over a clear policy horizon (2) a measure of uncertainty around the outlook, (3) specification of the interest rate path upon which the outlook is conditioned, and (4) the provision of forward guidance. Perhaps surprisingly, central banks that operate under similar inflation targeting frameworks — as established in Table 3.1 — do not necessarily employ similar communication strategies. For the sake of brevity, I refrain from evaluating each of the central banks in the sample and instead assess the communication strategies of select central banks to highlight common themes and key differences.

For the Bank of Canada, the policy horizon can be easily identified in its *Monetary Policy Report* (*MPR*) from July 2022. The inflation outlook extends over a forecast horizon of approximately 10 quarters, with a decomposition chart illustrating contributions to the expected deviation of inflation from target from various factors such as commodity prices, supply chain disruptions, and inflation expectations. These factors are referenced in support of the written outlook, in which the Bank expects "a decline in oil prices and lower house prices [to] pull down quarter-over-quarter inflation in the second half of 2022, [such that] inflation returns to the 2–percent target by the end of 2024 as global price pressures and domestic demand ease" (Bank of Canada, 2022c). Taken together, the policy horizon can be identified as 10 quarters, or 2 $\frac{1}{2}$ years.[1] However, the Bank does not provide any measure of uncertainty around its outlook and the underlying interest rate path is not specified. Forward guidance is provided in the press release accompanying the *MPR*, stating that

[1] A 10–quarter policy horizon is somewhat unusual, exceeding the typical horizon of 6 to 8 quarters stipulated in the Bank's monetary policy framework. The current trend of longer-than-average policy horizons extends to many of the central banks within this discussion, as developed economies across the world continue to experience elevated levels of inflation.

"the Governing Council continues to judge that interest rates will need to rise further, and the pace of increases will be guided by the Bank's ongoing assessment of the economy and inflation" (Bank of Canada, 2022a).

The European Central Bank (ECB) takes a very similar approach to the presentation of its inflation outlook, with contributions to inflation from various factors specified in detail through text alone rather than a decomposition chart. For example, the Eurosystem staff macroeconomic projections from June 2022 state that "HICP inflation is expected to start declining toward the end of 2022, mainly owing to downward base effects in the energy component and the assumed decline in the price of oil in line with futures prices. [...] Food inflation is also expected to start moderating in 2023, as the upward pressures from energy input costs, fertiliser prices and prices for food commodities subside." From this outlook, inflation returns to target at the end of 2024, implying a policy horizon of 10 quarters. Having recently exited a prolonged period of negative rates, the ECB's forward guidance simply reiterated that "further normalisation of interest rates will be appropriate" (European Central Bank, 2022).

In comparison, the US Federal Reserve's inflation outlook does not provide sufficient information to identify a range of policy horizons and relies upon a much less technical approach to forecasting. For its inflation forecast in the June 2022 *Monetary Policy Report*, the Fed reports the annualized median, central tendency, and range for inflation for 2022 to 2024 as projected by members of the Federal Reserve Board and the Federal Reserve bank presidents. As the median inflation forecast does not reach the 2–percent target by 2024, the Fed indicates that the "longer-run" median is anchored at 2 percent. In its narrative outlook, the Fed simply states that "the Committee is strongly committed to returning inflation to its 2 percent objective" (Board of Governors of the Federal Reserve System, 2022b). The median projection for inflation is contextualized with a 70–percent confidence interval based on historical forecast errors, and the participants' assessments of the uncertainty around their projections. A projection for the Fed funds rate is also reported, with each participant providing point forecasts for the "appropriate monetary policy" over the forecast horizon, which are aggregated into a median outlook for the policy rate. On forward guidance, however, the Fed takes the same vague approach as its inflation outlook, stating that it "anticipates that ongoing increases in the target range will be appropriate" Board of Governors of the Federal Reserve System (2022a). While the policy horizon cannot be reasonably identified from the information provided, this omission reflects the challenges posed to central bank communications by the extraordinary level of inflation with which the Fed must contend. Revisiting the February 2020 *Monetary Policy Report* — at the time of which economic conditions were much more stable — the median projection for inflation is exactly 2 percent for both 2021 and 2022. While the presentation of the forecast as an annual average precludes the policy horizon to be precisely identified, it is nevertheless evident that under normal circumstances, it is customary for the Fed's inflation outlook to converge toward target over the projection horizon.

The Swiss National Bank illustrates a different approach to conditioning an inflation outlook upon an interest rate path. In its 2022Q2 *Quarterly Bulletin*, the SNB presents two sets of conditional forecasts through 2025Q1. The first set — conducted in March 2022 — holds the SNB policy rate at 75 basis points over the full forecast horizon, while the second — conducted in June — holds the policy rate at 25 basis points. In general, the SNB produces its conditional forecasts by fixing the policy rate at its most recent value. Under the updated forecast from June, inflation is expected to fall below the 2–percent upper bound in 2023Q2, implying a policy horizon of 5 quarters. Comparatively, the policy horizon embedded in the March outlook is somewhat shorter at 4 quarters. However, neither forecast incorporates any measure of uncertainty. On forward guidance, the SNB states that "it cannot be ruled out that further increases in the [...] policy rate will be necessary in the foreseeable future to stabilise inflation in the range consistent with price stability over the medium term" (Swiss National Bank, 2022b).

Recall from the preceding section that the Bank of England is required to submit an open letter to the Chancellor whenever inflation deviates from the 2–percent target by more than 1 percentage point. Relative to other central banks in the sample, this additional emphasis on transparency is reflected in the Bank's presentation of its inflation outlook. In the August 2022 *Monetary Policy Report*, the inflation forecast is shown as a fan chart with 30–, 60–, and 90–percent confidence intervals. The forecast is conditioned upon "the paths for policy rates implied by financial markets," from which both the Bank rate and the effective exchange rate are mapped as annual averages over the 3–year projection horizon (Bank of England, 2022). Taking this as given, the Bank also reports an alternative outlook holding the Bank rate constant at 175 basis points. However, neither forecast shows inflation returning to target, implying that the policy horizon extends beyond 13 quarters but cannot be otherwise identified. Despite reiterating the Bank's obligation to address the appropriate policy horizon in an open letter to the Chancellor, the latest letter from June 2022 fails to do so convincingly. Specifically, the Bank states that "the MPC will take the actions necessary to return inflation to the 2–percent target sustainably in the medium term, in line with its remit. The scale, pace and timing of any further increases in the Bank Rate will reflect the Committee's assessment of the economic outlook and inflationary pressures" (Bailey, 2022). Similar to the Federal Reserve, the omission of a precise policy horizon appears to be the exception rather than the rule. The Bank of England's *Monetary Policy Report* from January 2020, for example, shows inflation converging to the 2–percent target by 2022Q1, implying a policy horizon of 8 quarters. However, in each of its five most recent open letters dating back to May 2021, the Bank falls short of identifying a policy horizon with greater specificity than "over the medium term". This suggests that under extenuating circumstances when an unusual degree of flexibility needs to be exercised, the Bank has struggled to meet the standard of transparency established in its mandate.

Similar to the Federal Reserve and the Bank of England, the Norges Bank's latest inflation outlook in its 2022Q2 *Monetary Policy Report* does not show inflation returning to target by the end of its conventional forecast horizon of 3 years. To address this gap and clarify its expectation for a return to "a normal situation," the Bank conducts an additional set of forecasts that extend 1 year beyond the forecast horizon. These forecasts are presented with the caveat that "the extended forecasts are derived directly from our macroeconomic model NEMO and therefore follow more mechanically from our model than the forecasts within the current horizon" (Norges Bank, 2022). With inflation returning to target by the end of 2027, the policy horizon is implied to be 23 quarters. Acknowledging the elevated degree of uncertainty associated with forecasting over such a long horizon, the Bank also provides extended forecasts for its policy rate and the output gap to support the inflation outlook. Alongside the forecasts, the Bank reports the actual monetary policy reaction function that is used to endogenously determine the policy rate within their macroeconomic model — the only institution to do so within the sample. However, despite the wealth of technical information, the inflation outlook is presented without a measure of uncertainty.

Taken together, this brief review of communication strategies illustrates that when policy frameworks are put into practice, the proposition of flexible inflation targeting is made more challenging by the need to balance credibility and the management of inflation expectations with the uncertainty inherent to a dynamic and globalized economic system. While far from comprehensive, the examples above highlight the key features that differentiate the strategies employed by OECD central banks. The varying information content and methods of presentation also reveal the relative priorities placed on varying aspects of central banks' respective mandates.

Some central banks, such as the Bank of England, provide extensive information to the public, augmenting their inflation outlook with confidence intervals, a conditional interest rate path, an alternative scenario, and a mechanism to affirm public trust through open letters when large shocks to inflation occur. However, as evident from the discussion above, this emphasis on transparency may at times place undue pressure on credibility. On the other hand, central banks such as the Federal Reserve are much more comfortable with open-ended, arguably imprecise communications. In the context of the Fed's updated framework of average inflation targeting, deviations from target may be more frequent and at times intentional. Without committing to explicit outcomes over a specified horizon, the Fed has the flexibility to balance its management of medium-term price stability with the policy space to pursue its dual mandate — as long as its credibility remains intact. Central banks such as the Bank of Canada and the ECB hold a middle ground in which the policy horizon is clearly communicated and forward guidance is used to establish expectations for the general trajectory of monetary policy, but measures of uncertainty and key assumptions are omitted. While this approach appears balanced in most circumstances, the one-dimensional nature of forecasts at times leaves information to be desired when the level of uncertainty in the economy becomes elevated.

As each approach has advantages and shortcomings under different economic circumstances, it is difficult to normatively assess the various communication strategies considered in this discussion. For a given central bank, the presentation of its inflation outlook is also dynamic — the extent and manner in which information is communicated evolves over time. For this reason, the following analysis of the Bank of Canada will explore the temporal dimension of its communication strategy in much greater detail while expanding upon the key elements of flexible inflation targeting established in this section.

3.2 The Bank of Canada — A Detailed Review

To motivate the empirical analysis in Chapter 4, I first embark upon an in-depth review of the Bank of Canada's inflation targeting framework and the evolution of its communication strategy across different leadership regimes. To match the time frame of the empirical analysis, I will review communications from the tenures of Governors David Dodge (2001M2–2008M1), Mark Carney (2008M2–2013M5), Stephen Poloz (2013M6–2020M6), and Tiff Macklem (2020M7–present).

Jointly with the Ministry of Finance, the Bank of Canada formally commits to inflation targeting in 1991. In the initial phase, the Bank is tasked with bringing inflation down from 6 percent in 1991M2 to 3 percent by the end of 1992, and subsequently to 2 percent by the end of 1995 (Bank of Canada, 1991). In 1993, the inflation-reduction program is extended to maintain inflation within a range of 1 to 3 percent — with a 2–percent midpoint — from 1995 to 1998. The same parameters for inflation control are extended again from 1998 to the end of 2001. In early 2001, the extension is formalized into a comprehensive document entitled the Renewal of the Inflation-Control Target.[2] Published in coordination with the Ministry of Finance every 5 years, this document provides background information on the inflation targeting regime, assesses Canada's experience under inflation targeting, motivates and explains any changes to the framework, and addresses potential avenues for further research. The *Renewal*, alongside the Bank's quarterly flagship publication — the *Monetary Policy Report* (*MPR*) — are the key resources upon which my analysis of the

[2]Hereafter referred to as the *Renewal*.

Bank's communication strategy will draw.

3.2.1 Dodge Era (2001M2–2008M1)

The first release of the *Renewal* in May 2001 — coinciding with the start of Governor David Dodge's tenure — establishes an inflation targeting framework from which considerations of flexibility are effectively absent. The inflation target does not incorporate complementary objectives such as financial stability or full employment, implying that the Bank is primarily concerned with returning inflation to the target midpoint as quickly as possible. As a result, the Bank views short-run deviations of inflation from target as consequences of the transmission lag of monetary policy and uncertainty around shocks, which are accommodated by the target range of ±1 percent. However, a small degree of variation is specified in the optimal policy horizon, for which the Bank determines that "monetary policy will therefore be directed to moving inflation to the target midpoint over a 6-to-8-quarter horizon" (Bank of Canada, 2001).

An additional feature of the framework established under the 2001 *Renewal* is the orientation of monetary policy around core inflation rather than total CPI inflation. A new preferred measure of core inflation, CPIX, is introduced, replacing the previously preferred measure of total CPI excluding food, energy, and the effect of changes in indirect taxes. CPIX excludes from total CPI "eight of the most volatile components identified by the Bank of Canada,"[3] and the effect of changes in indirect taxes. The Bank motivates its preference for core inflation as "a shorter-term operational guide to policy" as follows:

> "One reason for this focus on core inflation is the considerable short-run volatility in certain components of the CPI. Since movements in these volatile elements tend to reverse themselves fairly quickly, and since monetary policy actions affect inflation over a longer period, it would be inappropriate for monetary policy to try to offset the short-run movements in the total CPI caused by these fluctuations" (Bank of Canada, 2001).

Accordingly, the policy horizons associated with the Bank's inflation outlook — published biannually in the April and October *MPR*s — generally anchor CPIX inflation to the 2–percent target while at times allowing total CPI inflation to deviate at the end of the forecast horizon. In this context, from April 2001 to October 2006, the sequence of policy horizons is remarkably consistent. Despite multiple runs during which CPIX inflation approaches or exceeds the target band, the policy horizon for core inflation never breaches 8 quarters. There are, however, instances in which relatively minor deviations from the 2–percent target

[3] These are fruit, fruit preparations and nuts; vegetables and vegetable preparations; mortgage interest costs; natural gas; fuel oil and other fuels; gasoline; inter-city transportation; and tobacco products and smokers' supplies (Statistics Canada Table 18100256)

are projected to dissipate over a horizon shorter than 6 quarters. For example, with CPIX averaging 1.6 percent in 2005Q3, the Bank states in October 2005 that "the core rate is projected to rise to 2 percent by mid-2006 and remain there through 2007" — implying a policy horizon of 3 to 4 quarters. While the updated framework in the 2006 *Renewal* retains the typical horizon of 6 to 8 quarters, it adds the caveat that "specific occasions may arise in which a somewhat shorter or longer time horizon might be considered"(Bank of Canada, 2006).[4] Against the backdrop of strong domestic growth and stable inflation, the April and October *MPR*s in 2007 — the last two publications under Governor Dodge — report policy horizons of 3 and 4 quarters, respectively.

The transparency of the Bank's communication strategy over Governor Dodge's tenure is partially undermined by its presentation of inflation forecasts. Until the April 2003 *MPR*, the point forecasts for core and total CPI inflation are communicated solely through text at arbitrary intervals, lacking visual representation. For example, the inflation outlook from October 2002 reads "the overall effect of these one-off factors is expected to peak in the fourth quarter of this year, pushing core inflation to about 3 percent by year end. Through 2003, these effects are expected to diminish, so that in the second half of the year, core inflation should return to its 2 percent target" (Bank of Canada, 2002). While one could interpret "the second half" as either 2003Q3 or Q4 and infer that the policy horizon is 4 to 5 quarters, the precise path of inflation and timing for convergence are both unclear. The introduction of a summary table in the April 2003 *MPR* for the Bank's economic outlook alleviates this concern to some extent. However, the intervals over which the *MPR* reports point forecasts remain inconsistent — for the preceding quarter and 2 quarters ahead, it presents the forecast as a quarterly average of year-over-year percent change, and for the second year of the horizon, it presents forecasts 6–month averages. The April 2006 *MPR* adds an additional year to the forecast horizon as an annual average. Despite these changes, the inconsistent intervals used over the forecast horizon do little to address the imprecision of the communicated policy horizons.

Two key characteristics stand out from the first years of Governor Dodge's tenure — the reliance of monetary policy on core inflation, and the lack of flexibility in the inflation targeting framework. Although the 2006 *Renewal* does not formalize the concept of flexibility, the policy horizons embedded in the inflation outlook exhibit an inclination to depart from the optimal target horizon of 6 to 8 quarters. In the absence of significant economic shocks

[4] Around this time, the Bank begins to explore the implications of shorter and longer policy horizons. In particular, two in-house models, developed by Cayen et al. (2006) and Basant-Roi and Mendes (2007), are used to estimate optimal policy horizons. Following a "typical" macroeconomic shock, Cayen et al. (2006) estimate a range of 4 to 11 quarters under normal economic conditions. Incorporating financial shocks, Basant-Roi and Mendes (2007) estimate a range of 2 to 9 quarters. Despite the range of estimates suggesting a shorter optimal horizon on average, Coletti et al. (2006), in their summary of the two models, concludes that the Bank of Canada's "target horizon of 6 to 8 quarters remains appropriate in most instances."

and persistent inflationary pressure, it is at times natural for the policy horizon to fall short of 6 quarters. However, as the global economy slides into recession in 2008, the Bank's view of the flexibility at its disposal would shift quickly and dramatically.

3.2.2 Carney Era (2008M2–2013M5)

The first sequence of *MPR*s in Governor Mark Carney's tenure subtly but unmistakably signals a greater appetite for flexibility in the Bank's inflation targeting framework. With core inflation and total CPI inflation both below target at 1.4 percent and 1.8 percent respectively in 2008Q1, the policy horizon in the April 2008 *MPR* marks the first instance in which the target horizon may be exceeded. Specifically, with core inflation projected to increase to 2 percent in 2010, the policy horizon is 8 quarters if inflation returns to target in 2010Q1, and 11 quarters if it returns in 2008Q4 (Bank of Canada, 2008a). The October *MPR* subsequently confirms that convergence is expected by the end of 2010, leading to a policy horizon of 9 quarters (Bank of Canada, 2008b).

During the financial crisis from late 2008 to the end of 2009, policy horizons hover between 8 and 10 quarters. Over the same period, the Bank also begins to increasingly emphasize the risk posed by financial stability to its inflation outlook. For example, the July 2009 *MPR* states:

> "The principle downside risks to inflation relate mainly to the external sector. The restoration of normal financial conditions could be more gradual than expected and further setbacks cannot be precluded — triggered, for example, by unexpected losses at financial institutions or by escalating concerns about current account and budgetary imbalances. Such developments could have serious spillover effects in Canada through trade, financial, and confidence channels" (Bank of Canada, 2009).

In April 2009, the *MPR* itself also undergoes a significant transformation. Following a commissioned external review that compiled reactions from key users and analyzed general sentiments toward the content of the publication, the frequency of the *MPR* is increased from two to four publications per year — released quarterly in January, April, July, and October — accompanied by a visual redesign. While this redesign is primarily motivated by accessibility, the information content of the inflation outlook also dramatically improves. Beginning in the July *MPR*, the outlook for core and inflation CPI inflation shows quarterly point forecasts for the full 3–year forecast horizon rather than half-year and annual averages, identifying the policy horizons much more precisely. To illustrate, the outlook states that "core inflation is projected to return to 2 percent in the second quarter of 2011," leading to a policy horizon of exactly 8 quarters (Bank of Canada, 2009). In addition, the outlook for both measures also introduces line graphs of the base-case projection with 50– and

90–percent confidence intervals, roughly quantifying the uncertainty around the Bank's outlook.

As the global economy recovers from the financial crisis between mid-2009 and the end of 2011, the Bank signals an increased level of comfort with flexibility around its inflation target. Over this period, the policy horizons embedded in the inflation outlook span from 4 to 9 quarters. Reflecting upon the challenges of navigating monetary policy through the financial crisis, the Bank acknowledges in the 2011 *Renewal* that "the recent crisis has further refined [our] understanding of how [...] flexibility should be applied" (Bank of Canada, 2011). Specifically, "owing to the stability of Canada's financial system over the past two decades, the Bank has not actually had occasion to make pre-emptive use of the flexibility of the inflation targeting regime in response to a buildup of financial imbalances" (Bank of Canada, 2011). However, given its assessment of the post-crisis economic environment and the corresponding role of monetary policy, the Bank formally establishes the core motivation for the pursuit of flexible policy horizons. That is, "because the consequences of financial excesses may be felt over a longer and more uncertain horizon than other economic disturbances, the potential may exist for tension among output, inflation and financial considerations over the typical 2–year monetary policy horizon. In these circumstances, the Bank would need to use the flexibility available to it [. . .] to bring inflation back to target over a somewhat longer horizon" (Bank of Canada, 2011).

To this effect, the *Renewal* enumerates a list of broad circumstances under which it may be optimal to make use of flexibility. First, "unusually large and persistent shocks" may warrant a longer horizon to maintain economic and financial stability. Second, a longer horizon may allow monetary policy to "promote adjustments to financial excesses or credit crunches." Finally, the policy horizon may need to accommodate the degree of uncertainty surrounding shocks to the economy and minimize potential adverse consequences that may arise from the realization of downside or upside risks. The updated framework views the inflation target as flexible in that the objective of returning inflation to target needs to be balanced with the impact of monetary policy on financial stability. In the Bank's own words, "a framework anchored on a solid and credible inflation target provides the flexibility for monetary policy to play an occasional role in supporting financial stability" (Bank of Canada, 2011).

While the target horizon established in the 2001 *Renewal* is omitted from the 2011 *Renewal* altogether, it is clear that the Bank still intends, when feasible, to orient its policy actions around that range. Beginning in January 2012, each publication of the *MPR* reiterates in its introduction that "typically, the Bank seeks to return inflation to target over a horizon of 6 to 8 quarters. However, the most appropriate horizon for returning inflation to target will vary depending on the nature and persistence of the shock buffeting the economy"

(Bank of Canada, 2012). In this context, with inflation exhibiting a distinct downward trend toward the bottom of the target range, the policy horizons from 2012 to the end of Governor Carney's tenure in 2013M5 remain remarkably consistent, ranging from 7 to 9 quarters.

As the Bank navigates through the first recession since its adoption of an inflation target, Governor Carney's tenure is confronted with much higher levels of uncertainty and volatility than his predecessor's. Given the persistence and magnitude of shocks associated with the financial crisis, the Bank's policy actions under Governor Carney are characterized by longer policy horizons near the top of — and frequently exceeding — the 6-to-8–quarter target horizon. The culmination of these experiences gives rise to significant changes to the inflation targeting framework. Most notably, the acknowledgment of financial stability as a relevant consideration for monetary policy, along with other lessons from the financial crisis, lead the Bank to expand and formalize its approach to flexibility. In addition, improvements to the presentation of the *MPR* — particularly to the precision and information content of the inflation outlook — effectively bolster the transparency of policy communications.

3.2.3 Poloz Era (2013M6–2020M6)

Assuming office in June 2013, the economic conditions Governor Stephen Poloz inherited from his predecessor are considerably more favourable than the hand-off between Governors Dodge and Carney in 2008. However, in June 2014, global commodity prices begin a precipitous collapse that would weigh heavily on inflation for the subsequent two years. As core inflation begins to diverge from total CPI inflation and the outlook for the latter becomes increasingly uncertain, the *MPR* introduces additional information to anchor the inflation outlook. First, alongside fan charts, variation in private sector forecasts is introduced as an additional measure of uncertainty around the total CPI inflation. In the July 2014 *MPR*, this reads "based on the past dispersion of private sector forecasts, a reasonable range around the base-case projection for total CPI inflation is ± 0.3 percentage points" (Bank of Canada, 2014a). Second, alternative scenarios for inflation are added to the outlook discussion. For core inflation, alternative paths are conditioned upon different assumptions for the output gap; for total CPI inflation, they are conditioned upon different assumptions for oil prices. For example, the October 2014 *MPR* states that "if the base-case scenario assumed that oil prices were 10 percent higher (lower), total CPI inflation would be higher (lower) by 0.3 percentage points over the coming year" (Bank of Canada, 2014b). Finally, in the January 2015 *MPR*, a decomposition chart is added to the core inflation outlook, illustrating the percentage point contributions of various categories toward the deviation from target.[5] Taken together, Governor Carney's previous efforts toward greater transparency in the inflation outlook are reinforced in the first years under Governor Poloz. Between July 2014 and October 2015, the policy horizons for core inflation range from 4 to 9 quarters, shortening as the commodity price shock dissipates.

In the January 2016 *MPR*, the discussion around inflation shifts unexpectedly. While core inflation — CPIX specifically — has been the primary focus of the inflation outlook and the associated policy horizons since 2001, discussions suddenly become oriented around total CPI inflation. The order of presentation in the inflation outlook emphasizes total CPI over core inflation, and the decomposition chart shows contributions to the deviation of total CPI from target.[6] Alongside these changes, the Bank also adopts in the summary table of its inflation outlook "a convention of showing the last quarter of historical data and the next 3 quarters, [. . .] as well as fourth-quarter-over-fourth-quarter projections of [. . .] inflation for longer horizons" (Bank of Canada, 2016a). Somewhat similar to the original summary tables under Governor Dodge in which 6–month and annual averages are reported for the latter half of the projection horizon, this convention once again obscures the precise identification of policy horizons. Compounding upon this issue, the Bank also implements a subtle change to the tolerance threshold for convergence toward the inflation target. To illustrate, in the January 2016 *MPR*, the outlook states that "in 2017, [...] total CPI inflation and core inflation are projected to be close to 2 percent" (Bank of Canada, 2016a). While previous forecasts converge to exactly 2 percent, at the end of the projection horizon the forecasts for total CPI and CPIX inflation are both 1.9 percent, indicating an updated tolerance of ± 0.1 percentage points. Taking this as given, the policy horizon may be anywhere between 5 and 8 quarters. The October 2016 *MPR* projects total CPI inflation to settle at 1.9 percent in 2017Q4 and 2.0 percent in 2018Q4. Without clarifying the specific threshold for convergence and presenting forecasts for the intervening quarters, the policy horizon could be reasonably interpreted to be anywhere between 2 and 9 quarters.

Taken together, the changes implemented in January 2016 mark a significant step backward from the efforts made under Governor Carney toward greater transparency. However, there is little explanation or motivation to support this shift when the 2016 *Renewal* is published in November. The *Renewal* also does little to update or expand the Bank's approach to flexibility, simply reiterating that "the flexibility in the monetary policy framework allows the Bank to opt for a policy path that aims to return inflation to target over a longer time frame than normal so that its policy actions do not significantly worsen financial stability concerns" (Bank of Canada, 2016b). Instead, the change discussed at length is the replace-

[5] The initial categories were output gap and retail competition, exchange rate pass-through, and sector-specific factors.

[6] The new categories are commodity prices excluding pass-through, output gap, exchange rate pass-through, and others.

ment of CPIX as the Bank's preferred measure of core inflation. In its place, a set of three preferred measures — CPI-trim, CPI-median, and CPI-common — are to be monitored together as "multiple indicators will help the Bank transparently manage the risks associated with the shortcomings of any single indicator" (Bank of Canada, 2016b).

Between January 2017 and October 2019, inflation experiences little volatility and remains strictly within the target band. While this likely leads to relatively short policy horizons over this period, due to the changes discussed above, the precise range of horizons is difficult to identify. For example, the inflation outlook from the October 2019 *MPR* simply states that "the Bank forecasts CPI inflation will be around the 2 percent target over the projection horizon" — from which one could infer that the policy horizon falls between 5 and 9 quarters. Starting in the January 2017 *MPR*, the Bank further reduces the information content of the inflation outlook. The *MPR* removes the fan charts, relegates the 50– and 90–percent confidence intervals to the discussion, and omits the projections for core inflation altogether.

The onset of the Covid-19 pandemic in early 2020 presents the most significant challenge for Governor Poloz just months before the end of his 7–year term. In the April *MPR* — the last of Governor Poloz's tenure — the Bank makes the unprecedented choice to forego a base-case projection altogether. In place of the conventional summary table, the *MPR* presents the inflation outlook as a wide range of scenarios that bounds inflation between approximately -0.5 and 2.5 percent in the near term and converges to ± 0.25 percentage points around the 2–percent target by the end of 2022. Even with the extraordinary breadth of potential outcomes, the Bank nevertheless specifies that "this range understates the uncertainty in the outlook for inflation because there are both upside and downside risks to each of these projections" (Bank of Canada, 2020). By the time Governor Poloz departs in June, total CPI has recovered from a brief bout of deflation in the two preceding months but would remain below the target range for the rest of the year.

Setting aside the first months of the pandemic, the evolution of the Bank's inflation targeting framework over Governor Poloz's tenure is primarily defined by changes to the communication strategy of the inflation outlook in the *MPR*. From the perspective of transparency, these changes arguably unwind the incremental progress gained under Governor Carney and effectually undermine the public's ability to identify the policy horizons embedded in the outlook. Additionally, while the role of flexibility does not materially change over this period, the reduced emphasis on core inflation in favour of total CPI inflation has been retained as a defining feature of the inflation targeting framework under which the Bank operates today.

3.2.4 Macklem Era (2020M7–)

Upon Tiff Macklem's return to the Bank as Governor, his experience navigating the financial crisis at the Department of Finance is immediately called upon to address the initial impacts of the Covid-19 pandemic. By July 2020, although the Bank has already lowered the overnight rate to the effective lower bound, inflation has nevertheless fallen close to zero. In the July *MPR*, the Bank partially returns to its standard practice by presenting a *central* scenario, which relies more heavily on judgment than the conventional base-case projection and omits the uncertainty bands around the outlook. The Bank also refrains from showing a return to target within the projection horizon until the January 2021 *MPR*, which states that "although the timing is highly uncertain, inflation is expected to return sustainably to the 2 percent target as excess capacity is absorbed. In the base-case projection, this occurs in 2023" (Bank of Canada, 2021b). The policy horizon implied by this outlook is 9 to 12 quarters.

Due to a gradual recovery in demand and easing restrictions in some provinces, inflation rebounds much more quickly than expected over the second half of 2020, rising from 0.1 percent in July to 1 percent in November. Amplified by base-year effects from 12 months prior, inflation returns to 2 percent in March 2021, nearly 2 years earlier than projected. However, following the subsequent reading of 3.4 percent in April 2021, the Bank's focus quickly shifts to the upside risks to the outlook as the threat of pent-up inflationary pressure from a combination of factors begins to materialize.

The Bank emphasizes in the April 2021 *MPR* its view that the factors pushing inflation above target were temporary, stating that "the transitory increase reflects the fact that the prices of some goods and services fell sharply with the sudden drop in demand at the onset of the pandemic. This results in base-year effects — CPI inflation appears higher on a year-over-year basis because prices of gasoline, motor vehicles and traveller accommodations fell in early 2020. [...] In setting monetary policy, the Bank looks through such temporary movements in inflation"(Bank of Canada, 2021c). However, as restrictions ease in Canada over the summer months, global supply chain disruptions and a sudden increase in demand place additional pressure on inflation, and temporary factors associated with the pandemic prove to be more persistent than expected. Although inflation rises above 3.5 percent by July, the Bank nevertheless pledges to hold "the policy interest rate at the effective lower bound until economic slack is absorbed so that the 2 percent inflation target is sustainably achieved" (Bank of Canada, 2021a). As a result, the dynamics of the Bank's inflation outlook rely heavily upon the effects of these temporary factors dissipating without monetary policy intervention. Under this assumption, the Bank states in the October *MPR* that "CPI inflation is expected to ease from about 4.75 percent at the end of 2021 to around 2 percent at the end of 2022" (Bank of Canada, 2021d) However, given its "exceptional forward guidance," the Bank then expects inflation to bounce off of the target and moderately increase over 2023, returning toward 2 percent in 2024. In this scenario, the outlook implies policy horizons of 10 to 13 quarters.

As inflation continues to rise through the end of 2021, further upward pressure materializes in February 2022 as Russia's invasion of Ukraine sends global energy and food prices into turmoil and exacerbates global supply chain disruptions. With inflation above 5 percent, the Bank begins to raise its policy rate in March. Weighing the effects of energy prices, supply disruptions, and monetary tightening, the outlook in the April 2022 *MPR* sees inflation returning to target by the end of 2024. This implies a policy horizon of 11 quarters. In its discussion of risks around the outlook, the Bank also signals its concern that inflation expectations may become unanchored. Specifically, "higher inflation expectations could, in turn, lead to more pervasive increases in labour costs and inflationary pressures and could become embedded in ongoing inflation" (Bank of Canada, 2022b). With inflation rising to a 40-year high at over 7 percent in July, the Bank raises its policy rate by 100 basis points — the largest hike since August 1998 — for a total increase of 225 basis points since March. Conditional on this sequence of hikes and further tightening in the near term, the July *MPR* forecasts a return to target in 2024Q4, implying a policy horizon of 10 quarters. The orientation of the outlook around 2024Q4 in each of the three recent *MPR*s may indicate an effort to return to the 6-to-8–quarter target horizon, which in turn may dampen concerns around de-anchoring inflation expectations. However, as inflation continues to rise in the near term, the extent to which the Bank will be able to manage inflation expectations and navigate a soft return to target remains to be seen.

The 2021 *Renewal*, published in December, makes a concerted effort to distance its evaluation of the mandate from the ongoing challenges to inflation. The Bank conducts a full review of key alternatives to inflation targeting and places a particular emphasis on both inflation and maximum sustainable employment. While it falls short of committing to a dual mandate, the Bank nevertheless acknowledges that "full employment and output at potential are necessary conditions for achieving the inflation target on a sustained basis" (Bank of Canada, 2021e). To this effect, an additional dimension of flexibility is introduced to the inflation targeting framework such that "when circumstances warrant, the Bank will refrain from raising rates pre-emptively as inflation is approaching the 2 percent target. This will help the Bank better assess the current level of maximum sustainable employment" (Bank of Canada, 2021e).

Reflecting on its experience over the past 2 years, the Bank states:

> "During the pandemic, flexibility was exercised through an aggressive monetary policy response. The Bank quickly reduced the policy rate to its ELB, pro-

> vided forward guidance and used additional tools such as quantitative easing to support financial market functioning and reinforce the forward guidance. The forward guidance demonstrated that even without room to cut the policy rate further, the Bank can still provide monetary stimulus by making a commitment to hold the policy rate at its ELB for longer than the degree of excess capacity might suggest" (Bank of Canada, 2021e).

In retrospect, the Bank's commitment to its exceptional forward guidance and keeping the policy rate at the effective lower bound through 2021 may have exacerbated the inflationary pressures taking hold at the time of writing in 2022. Nevertheless, the *Renewal* affirms the Bank's appetite to expand its toolset in response to extraordinary circumstances. The use of forward guidance, unconventional monetary policy tools, and considerations of maximum sustainable employment represent extensions of the Bank's approach to flexibility beyond the lengths of policy horizons. The necessity of these extensions — both to achieving the inflation target and anchoring inflation expectations — has become distinctly evident over Governor Macklem's tenure, during which policy horizons have frequently exceeded the forecast horizon altogether.

The conduct of monetary policy since early 2020 has been particularly difficult to assess, as "extraordinary" economic circumstances have arguably encompassed all of the ensuing period up to the time of writing. In response to these extraordinary circumstances, Governor Macklem's tenure thus far may be characterized by the Bank's willingness — perhaps out of necessity — to fully utilize the flexibility at its disposal. However, beyond its management of the ongoing inflation crisis, the Bank must now also contend with the associated erosion of public trust and the escalation of political scrutiny around its independence. As it acknowledges in the 2016 *Renewal*, "the Bank's scope to exercise appropriate flexibility with respect to the inflation-targeting horizon is founded on the credibility built up through its demonstrated success in achieving the inflation target" (Bank of Canada, 2016b) With this credibility now under threat, the legacy of Governor Macklem's tenure will be defined by his ability to restore it.

3.2.5 Key Findings

The discussion above illustrates that the Bank of Canada's inflation targeting framework — and its approach to flexibility — has evolved considerably since 2001. From the initial commitment to a 6-to-8–quarter target horizon, economic circumstances and a refined understanding of the framework have led to a greater appetite for flexibility. The range of policy horizons, falling within or below target under Governor Dodge, is extended beyond 8 quarters under Governor Carney to accommodate the shock of the financial crisis. In the 2011 *Renewal*, this experience motivates the addition of financial stability as an explicit consideration for monetary policy, normalizing the pursuit of policy horizons beyond 8 quarters. The unprecedented magnitude of the Covid-19 crisis pushes policy horizons further, at times beyond the scope of the Bank's projection horizon, and illustrates the need for additional mechanisms to support the inflation target. As a result, the 2021 *Renewal* expands the basis for which flexibility may be utilized, including forward guidance, quantitative easing and tightening, and considerations of maximum sustainable employment.

The evolution of the Bank's communication strategy in the *MPR* has been much less linear. Significant strides are taken under Governor Carney to elevate the level of transparency exercised in the presentation of the inflation outlook and the discussion of policy horizons. While this push toward transparency initially continues under Governor Poloz's leadership, the trajectory leading up to the 2016 *Renewal* instead shows a preference toward the wider margin of error granted by comparatively more vague communications. As the Covid-19 pandemic unfolds, the elevated degree of volatility emphasizes the vital role of transparency in the management of public trust and inflation expectations. To this effect, even as the outlook itself remains uncertain, the Bank makes efforts to accurately communicate the assumptions and motivations behind its policy choices. However, the impact and timing of those choices now find themselves under an unprecedented amount of scrutiny as the Bank attempts to rein in a level of inflation not seen since the inception of inflation targeting in Canada.

Just as the expansion of flexibility in the 2011 *Renewal* is shaped by lessons learned from the financial crisis, the upcoming *Renewal* in 2026 will no doubt reflect upon experiences from the first years of Governor Macklem's tenure. Until then, with the current inflation crisis showing few signs of resolution, the direction toward which the role of flexibility will tend remains uncertain. Having reviewed the evolution of the Bank of Canada's inflation targeting framework and its approach to flexibility thus far, I now turn to the empirical evidence.

Chapter 4

Empirical Analysis

4.1 Theoretical Framework

4.1.1 A Simple Fisherian Model of Inflation

The empirical methods used in this study build on Mu and Voss (forthcoming), which is a work in progress on a larger, more comprehensive empirical analysis of monetary policy horizons. I will focus on a univariate autoregressive model of Canadian inflation, constructing a sequence of conditional forecasts over a significant portion of the Bank of Canada's inflation targeting period. I then examine the number of periods it takes for these forecasts to converge to the Bank's 2–percent inflation target, estimating the policy horizon at each point in time over the forecast horizon. The behaviour of the sequence of estimated policy horizons provides information about the operation of monetary policy in Canada over the sample period and the extent to which monetary policy is flexible under a very specific interpretation. In the following discussion, I present a simple model of monetary policy and inflation from Davig and Foerster (2021) — with minor changes — to motivate the empirical analysis.

The model is comprised of three equations. First, the standard Fisher equation is, in its log-linear approximation form:

$$i_t = E_t \pi_{t+1} + r_t \tag{4.1}$$

where i_t is the one-period nominal interest rate, π_{t+1} is the rate of inflation between period t and $t+1$, r_t is the *ex ante* one-period real interest rate, and E_t is the time t conditional expectation operator. The real interest rate is exogenous and follows an AR(1) process:

$$r_t = \rho r_{t-1} + \varepsilon_t \tag{4.2}$$

with $0 \leq \rho < 1$ and ε_t has a mean of zero and a variance of σ_ϵ^2. Finally, monetary policy follows a simple Taylor-type policy rule:[1]

$$i_t = \alpha \pi_t \tag{4.3}$$

The parameter α governs the monetary policy response to current period inflation and, following the Taylor principle, $\alpha > 1$. Also implicit in this policy rule is an inflation target of zero — $\pi^* = 0$. Combining the three equations, I obtain the non-explosive rational expectations solution for inflation,

$$\pi_t = \left(\frac{1}{\alpha - \rho} \right) r_t \tag{4.4}$$

This solution can also be represented as an AR(1) process:

$$\pi_t = \rho \pi_{t-1} + \left(\frac{1}{\alpha - \rho} \right) \varepsilon_t \tag{4.5}$$

$$\pi_t = \rho \pi_{t-1} + e_t \tag{4.6}$$

where $e_t = 1/(\alpha - \rho)\varepsilon_t$. The innovation e_t has mean zero and variance $\sigma_e^2 = \sigma_\varepsilon^2/(\alpha - \rho)^2$. The unconditional mean for inflation is zero and the unconditional variance for inflation is:

$$\sigma_\pi^2 = \left(\frac{1}{\alpha - \rho} \right)^2 \left(\frac{1}{1 - \rho} \right)^2 \sigma_\varepsilon^2 \tag{4.7}$$

More succinctly, this can be written as $\sigma_\pi^2 = \sigma_\pi^2(\alpha, \rho, \sigma_\varepsilon^2)$. The associated conditional forecasts at time t are:

$$E_t \pi_{t+j} = \rho^j \pi_t \qquad \forall j \geq 1 \tag{4.8}$$

Since π_t is a stationary AR(1) process, its conditional forecasts converge monotonically. However, as emphasized by Davig and Foerster (2021), the horizon for inflation to return to exactly zero is infinite. Making the policy horizon meaningful requires specifying a tolerance band, which I denote μ. Assuming $\pi_t > 0$, the policy horizon N can be defined as the number of periods it takes for the conditional forecast of inflation to converge to μ:

$$E_t \pi_{t+N} = \mu \tag{4.9}$$

[1] Davig and Foerster (2021) include an i.i.d. monetary policy shock in their treatment but this is not needed for this discussion.

By substituting Eq. (4.8), the policy horizon N_t can be identified as:

$$\rho^{N_t}\pi_t = \mu \tag{4.10}$$

I will revisit the properties of this condition in the context of flexibility in the upcoming section. Eq. (4.9) assumes that the process for inflation converges monotonically toward the inflation target (the unconditional mean) and that the inflation target is zero. More generally, the policy horizon N_t can be defined for a non-monotonic process for inflation and any inflation target π^* as follows:

$$|E_t\pi_{t+j} - \pi^*| < \mu \qquad \forall j \geq N_t(\rho, \pi_t; \mu) \tag{4.11}$$

The choice of μ is likely to be very influential over the lengths of policy horizons. One approach is to select a reasonable value for μ and estimate the range of policy horizons using conditional forecasts. One can then summarize an individual policy horizon as "given its current level, inflation is expected to return to within 25 basis points of target in 12 months".

Alternatively, building upon a suggestion in Davig and Foerster (2021) and Batini and Nelson (2001b), tolerance can be framed relative to the initial deviation of inflation from target. In other words, convergence occurs when χ percent of the initial deviation from target at π_t has permanently dissipated from the conditional forecast; or equivalently, when inflation has returned to within $1 - \chi$ percent of the initial deviation from target:

$$\mu(x, t) \equiv (1 - \chi)|(\pi_t - \pi^*)| \tag{4.12}$$

With this definition of tolerance, the condition Eq. (4.11) becomes:

$$|E_t\pi_{t+j} - \pi^*| < \mu \qquad \forall j \geq N(\rho, \pi_t; \mu(\chi, t)) \tag{4.13}$$

In both Eq. (4.11) and Eq. (4.13), the policy horizon is conditional on the parameters and the information set for the conditional forecast. For this model, these are ρ and π_t, respectively. More general models will have a richer set of determinants, as I discuss below.

By way of a concrete example, Figure 4.1 identifies both an absolute policy horizon (for $\mu = 0.25$) and a relative policy horizon (for $\chi = 0.90$) for a simple AR(2) model. Because this model has non-monotonic forecasts, it better demonstrates the determination of the policy horizons.[2] From the figure, $N = 10$ for $\mu = 0.25$; $N = 13$ for $\chi = 0.90$.

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

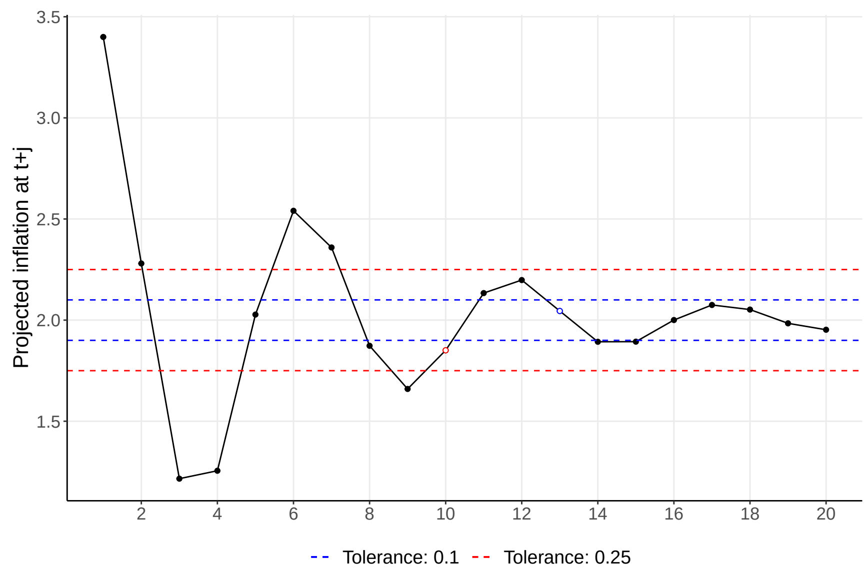

Figure 4.1: Conditional Inflation Forecast and Tolerance Bands

4.1.2 Applications of Flexibility

To illustrate the role of flexibility within the simple framework presented above, it is useful to revisit and explore the properties of Eq. (4.10). To recall, the condition for a positive π_t to converge to μ — given an inflation target of zero — is:

$$\rho^{N_t}\pi_t = \mu \tag{4.14}$$

This can be used to show that for a starting point where $\pi_t > \mu > 0$ and $0 < \rho < 1$, the policy horizon is longer for larger values of π_t — or more generally, for larger deviations of inflation from target.[3] While N_t does not directly depend on the monetary policy parameter α, this result reveals a type of flexibility that is nevertheless influenced by how monetary policy is implemented. To illustrate, consider an alternative policy rule where the central bank sets $i_t = r_t$ for all t. The Fisher equation collapses to $E_t\pi_{t+1} = 0$ and a stable solution for inflation is $\pi_t = 0$ for all t. In this case, monetary policy perfectly offsets the real interest rate shock in each period and the policy horizon is always zero. In comparison, under the original rule of $i_t = \alpha\pi_t$, the parameter α ensures that inflation does not immediately return to target; more generally, the policy horizon is flexible in that it depends upon the current level of inflation and how far away it is from target. While seemingly narrow in scope, in a broader context this is an important measure of flexibility for monetary policy. It says, very generally, as inflation is driven further from target the central bank conducts its monetary policy to return to target in a measured manner, rather than committing to a fixed horizon.

[3] For this discussion, I am assuming an absolute tolerance parameter μ consistent with Eq. (4.10).

This type of flexibility is the focus of my empirical analysis.

To get a better sense of how individual policy horizons vary with the central bank's choice of α, consider the distribution of N_t over some population $\mathbf{N}$. Although any individual N_t does not depend directly upon the policy parameter — which is α in the present case — the policy parameter does play a role in the distribution of N_t. In this model, I can conjecture that the variance of N_t will depend positively on the variance of π_t, which is previously specified as $\sigma_\pi^2(\alpha, \rho, \sigma_\varepsilon^2)$. From Eq. (4.7), an increase in α lowers the variance of π_t and hence the variance (range) of N_t. In other words, a more aggressive policy path keeps inflation closer to target on average and is therefore associated with a narrower range of policy horizons. Note, however, that ρ and σ_ε^2 will also affect the variance of N_t; specifically, increases in both parameters will likely increase the variation of π_t and therefore lead to a wider range of policy horizons. Furthermore, with a wider range, one can expect the mean and median values of policy horizons to increase, as these depend upon the distance — in absolute value — of inflation from target.

Given a full sample of π_t over T periods — which is used to estimate the sample counterparts of the population variables defined above — the parameters ρ, α, and σ_ε^2 can be assumed to be constant. The distribution of N_t is therefore determined by the central bank's choice of α and the properties of inflation, and any individual variation in N_t is due to variation in π_t only. More specifically, the length of the policy horizon at each time t is determined by idiosyncratic shocks to inflation drawn from the distribution of ε_t and its interactions with the parameters. If, however, the full sample over T periods were to be split into shorter sub-samples and estimated individually, then ρ, α, and σ_ε^2 would each be allowed to vary both between the split samples. Differences in the distributions of N_t across samples can therefore be attributed to any of the parameters and their interactions with each other, not just changes in α. The value of α may also adjust to changes in ρ or σ_ε^2 — or vice versa — further complicating the relationship between the parameters and the distribution of estimated policy horizons.

The above discussion demonstrates that for a given policy rule — represented by α — and an underlying economic structure — represented in this case by ρ and σ_ϵ — a path for inflation and the associated policy horizons N_t can be obtained. If any underlying parameters change, then the distribution of the policy horizons will also change. This result will be important for my empirical analysis, where one might expect changes in the underlying environment to have affected the estimated policy horizons.

4.2 Data

The focus of this chapter 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 and the Department of Finance in late February 1991 and initially stipulated a downward path for inflation from a rate of roughly 6 percent to 2 percent by 1995 (Bank of Canada, 1991). A target or target path beyond 1995 was not specified save for a general statement about further reductions to achieve price stability, thought to be less than 2 percent. By December 1993, inflation had fallen to 2 percent and the Bank of Canada and the Department of Finance agreed to a framework going forward of a 2 percent mid-point target within a 1-to-3–percent target range (Bank of Canada, 1993). This framework was to be in place until 1998 when it would be reviewed. As discussed in Section 3.2, this framework has been renewed numerous times with the same target and range, most recently in December 2021. For simplicity's sake, I consider 1991M1 to 2022M4 as the inflation targeting period. As I will address in Section 4.4, estimation of the empirical model for inflation is conditional on the 13 observations immediately preceding this period. Therefore, fixing $t = 1$ at 1991M1, I extend the sample back to 1989M12 ($N = 389$).

The Bank of Canada targets the year-over-year rate of change in the All-Items Consumer Price Index (CPI) as reported monthly by Statistics Canada, and this measure is the focus of this chapter.[4] The Bank has also referenced various *core* measures of inflation over the years as a guide for underlying inflation. The intention of these measures is to remove components that experience short-term — or *transitory* — volatility from All-Items CPI inflation to provide a better guide for forward-looking monetary policy.[5] Given the important role core inflation has had in Canadian monetary policy, I extend the empirical analysis to one of these core measures, CPIX. While it is not currently the preferred measure of core inflation, I view CPIX as a reasonable compromise as it served as the Bank's preferred measure of core inflation from 2001 to 2017, and is closely related to the measure used before 2001 (Macklem, 2001). It also offers the advantage of a complete time series, whereas the currently preferred measures of core inflation need to be spliced with predecessors that are not methodologically consistent.

For this study, I construct the rate of inflation as the year-on-year log differences of the underlying index, which is an approximation to the published rates of inflation targeted by

[4] In the previous chapter, the Bank of Canada refers to the reference price index for its inflation target as "Total CPI." Statistics Canada reports this series as "All-Items CPI." The two names are interchangeable, but I will defer to the latter to remain consistent with the data source.

[5] See Macklem (2001) and Khan et al. (2015) for earlier 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 currently preferred measures of core inflation.

the Bank of Canada. That is,

$$\pi_t = 100 \times (\ln P_t - \ln P_{t-12}) \tag{4.15}$$

where P_t is All-Items CPI or CPIX in levels. Neither series is seasonally adjusted.

Table 4.1: Summary Statistics

Sample : 1989M12-2022M4	Obs.	Mean	Median	Std. Dev.	Min	Max
Total CPI	389	2.0433	1.8714	1.2779	-0.9545	6.6818
CPIX	389	1.8958	1.7714	0.6948	0.6689	5.9437

Variable definitions:
(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].

Table 4.1 reports summary statistics for All-Items CPI and CPIX inflation from 1989M12 to 2022M4. Figure 4.2 plots All-Items CPI inflation in panel (a) and CPIX inflation in panel (b), each accompanied by the 2–percent target and the 1-to-3–percent target band. It is evident that while both measures of inflation are well-anchored around the 2–percent target, CPIX has generally exhibited less volatility. The latter result is expected as, by construction, measures of core inflation are intended to strip out short-term volatility to reveal the underlying — more stable — process of inflation over the medium term.

For both measures, peak inflation is observed at the start of the sample period when the Bank begins to implement its inflation targeting program, and at the end of the sample period as part of the ongoing experience of high inflation in Canada and beyond.[6] Meanwhile, troughs can be mapped to brief periods of deflation for All-Items CPI inflation in the mid-1990s, the initial recovery period from the financial crisis in 2009, and the beginning of the Covid-19 pandemic in 2020. While similar declines are observed for CPIX inflation over the same periods, it never experiences deflation and rarely falls below 1 percent. By and large, the Bank's inflation targeting framework has been, until recently, successful at maintaining low and stable inflation. I assume the time series are stationary over the full sample since a credible and constant inflation target necessarily implies that inflation is mean-reverting.

The underlying price indices, All-Items CPI and CPIX, exhibit considerable seasonality that

[6] As of 2022M4, inflation was slightly below the previous peak in 1991M1. However, in the months following the end of the sample period, inflation has well surpassed that peak.

Figure 4.2: Canadian Inflation (1989M12-2022M4)

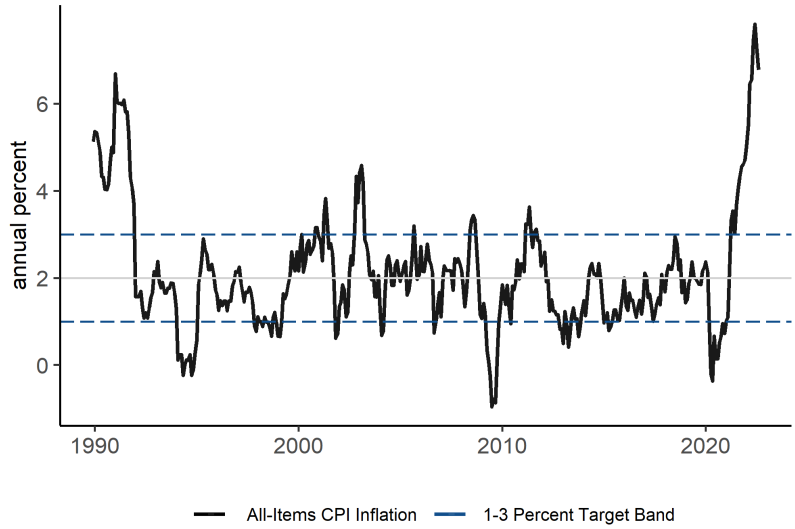

(a) All-Items CPI Inflation

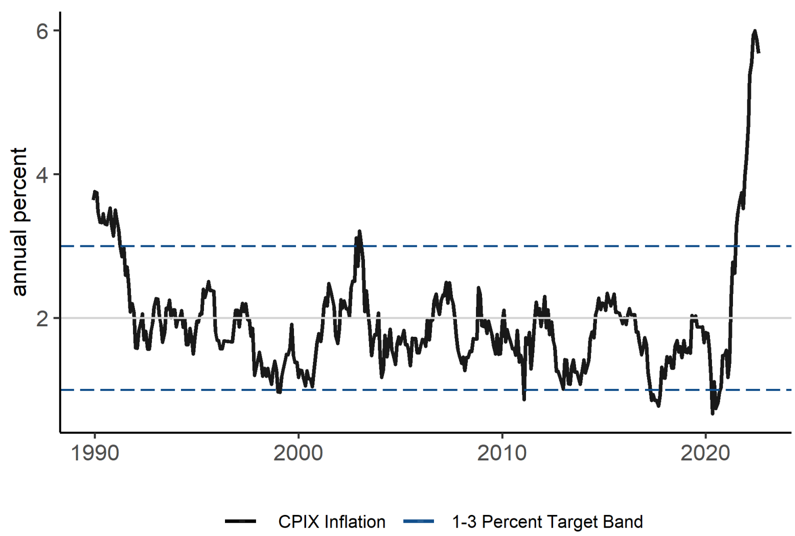

(b) CPIX Inflation

is, to some extent, controlled for by the seasonal differencing used to construct the inflation measures. The resultant inflation series, however, continue to exhibit a great deal of persistence at seasonal and near-seasonal frequencies, as evident in the partial autocorrelation functions in Figure 4.3. I address the issue of autocorrelation in Section 4.4.

Figure 4.3: Partial Autocorrelation Functions

(a) Total CPI Inflation PACF

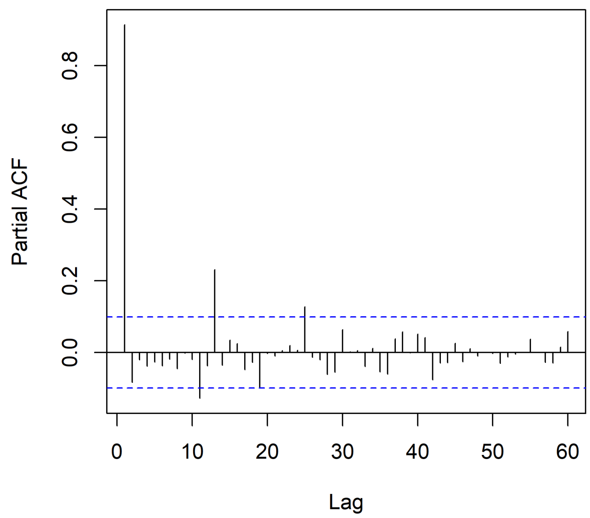

(b) CPIX Inflation PACF

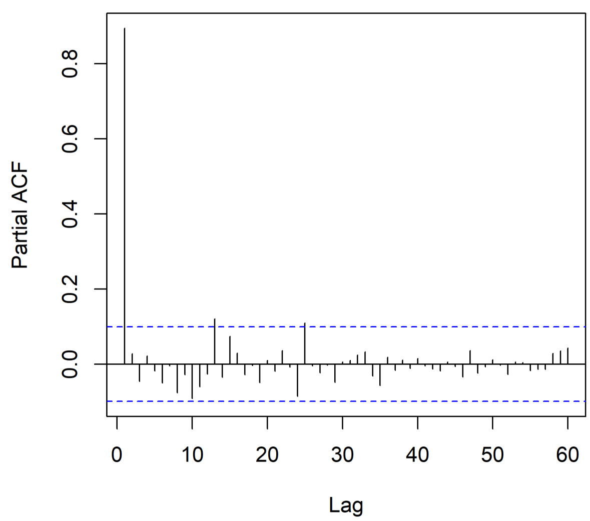

*Dotted lines denote 95% confidence interval.

4.3 Full Sample, Recursive, and Split Sample Estimation

One implication of the simple Fisher model of inflation from Section 4.1 is that the structure of the underlying economy, including central bank behaviour, determines the process for inflation. Recall from Eq. (4.5) that the solution for inflation in the model is an AR(1) process dependent upon the parameter ρ, arising from the underlying dynamics of the economy, and the parameter α, arising from the monetary policy rule. Under the strong assumption that the economic structure is stable and central bank behaviour is consistent, it might be reasonable to use the full sample of data to estimate the empirical model for inflation. For the AR(13) model used for my analysis, this implies that the estimated parameters characterize the *average* inflation process over the sample period. One potential disadvantage to this approach is that the policy horizons are estimated from within-sample forecasts. As a result, future information — not available to any central bank in practice — partially determines the dynamics of forecasted inflation.

One alternative approach is to estimate the process for inflation recursively. This approach estimates the model using an initial sub-sample, to which one additional observation is appended and the model is re-estimated, eventually converging to the full sample model. This approach produces out-of-sample forecasts for inflation, as the model only captures inflation's dynamics up to the last observation of a given sub-sample. Unlike the full sample approach, recursive estimation assumes that the model parameters ρ and α may vary over time, reflecting potential shifts in the macroeconomic structure and central bank behaviour. However, the dynamics associated with each additional observation are incorporated into the average inflation process up to that point in time, which does not capture structural shifts perfectly. One shortcoming of the recursive approach is that it requires the model to be reasonably well-behaved for each sub-sample. One could use a lag selection criterion to identify a unique lag length at each increment, but this becomes inefficient if the sample is large. I discuss this issue in greater detail in Section 4.4.

A more intuitive alternative to the recursive approach is split sample estimation, which divides the full sample into distinct periods that may correspond to unique identifying features. For example, the inflation data can be split into the time before, during, and after the 2008–09 financial crisis; or, it can be divided each time the Bank of Canada updates its *Renewal of the Inflation-Control Target*. To mirror the discussion in Section 3.2, the split samples used in the empirical analysis correspond to the completed tenures of the Bank's three most recent Governors.[7] A unique model is estimated for each split sample, with the associated policy horizons again drawn from within-sample forecasts as in the full sample approach. Contrary to the full sample, however, the forecasts constructed from the

[7]Govenor Macklem is excluded as he has served for less than 2 $\frac{1}{2}$ years at the time of writing.

split sample models do not depend on prior or future behaviour of inflation. The models can avoid estimation issues arising from potential structural shifts in the data if the split samples are well-chosen. However, the models may also behave poorly if the chosen sample periods are too short.

4.4 Lag Selection

Ideally, the statistical model should capture all of the relevant dynamics in the data. However, the discussion at the end of Section 4.2 suggests that this may be difficult to do in a parsimonious manner. Over the full sample, lag selection using a likelihood ratio (LR) test or a range of information criteria yield different results.[8] Given the limited sample size and the propensity for overfitting in higher-order autoregressive models, 60 lags (5 years) is assumed to be sufficient as the maximum lag length considered for capturing the dynamics of both series. The LR test proposes 49 lags for All-Items CPI with a p-value of 0.012 and 59 lags for CPIX with a p-value of 0.005, while the Akaike Information Criterion proposes 37 and 25 lags, respectively. The Schwarz Criterion proposes the fewest lags at 13 for both series.

Figure 4.3 shows that both time series only exhibit statistically significant autocorrelation up to and including the 13$^{\text{th}}$ lag at the 5 percent level, supporting the choice of the parsimonious specification. However, given the highly persistent nature of inflation data, the residuals from the AR(13) models retain some degree of autocorrelation when estimated over the full sample, as evident in Figure 4.4. This residual autocorrelation qualifies the key results discussed below in Section 4.6. I consider models with longer lags as a robustness check in Section 4.6.4.

For recursive estimation, the issue of lag selection is compounded by the need to choose a lag length that would be appropriate for each sub-sample. Setting the initial estimation sample to 1989M12 to 2000M1 and increasing the sample size by 1–month increments, there are 269 unique models that need to be estimated. Clearly, it would be inefficient to identify an appropriate lag length across multiple selection criteria for each model. A reasonable compromise, then, is to apply a common lag of 13 across all 269 models — consistent with the full sample estimation. While there may be some sub-samples for which the resulting model behaves poorly, this is true for any common lag length. By using a lag length that has been more methodically selected for full sample estimation, the forecasts constructed with the recursive estimates become increasingly more stable as they approach the full sample.

[8] The log-likelihood of a VAR with lag order p is tested against one with order $p\text{-}1$, starting with a maximum lag of ρ.

Figure 4.4: Partial Autocorrelation Functions of Model Residuals

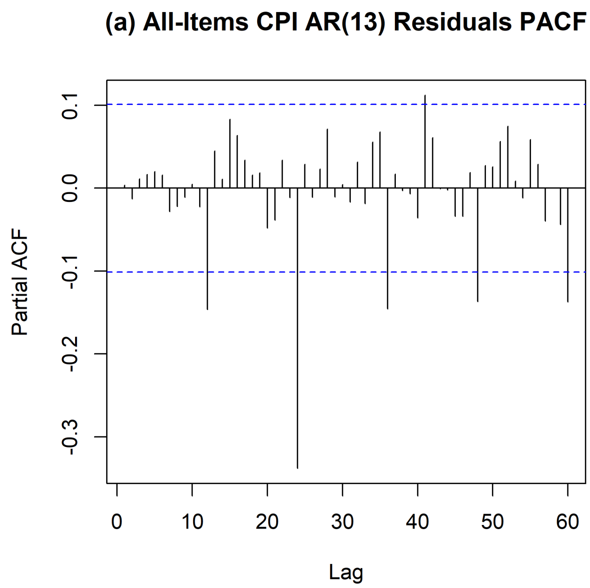

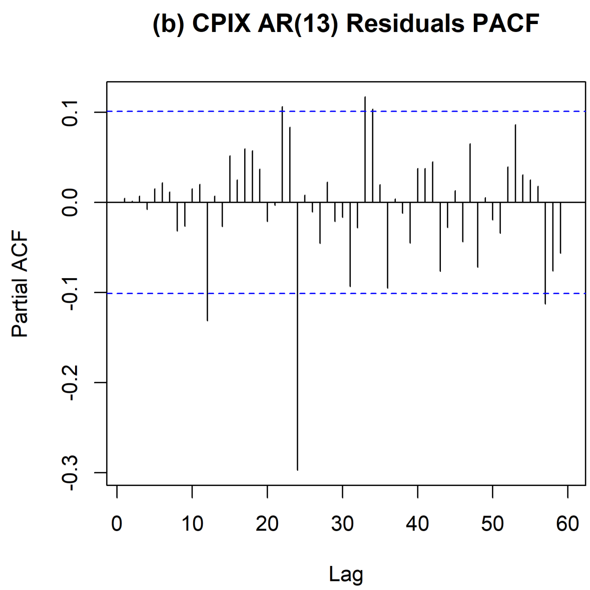

*Dotted lines denote 95% confidence interval.

Similarly, following the Schwarz Information Criterion, a lag length of 13 remains appropriate for estimating each of the split sample models associated with the tenures of recent Bank of Canada Governors. While other selection criteria suggest fewer lags, lag lengths longer than 12 months are needed to capture the seasonality inherent to inflation when calculated as a year-over-year change.

4.5 Tolerance Level and Convergence

Recall the convergence criteria set out in Eqs. (4.11) and (4.13). For a given tolerance level — either absolute or relative — and an inflation target π^*, a policy horizon can be calculated for each of the 269 sequences of conditional forecasts. Within this sample of policy horizons, I focus on the subset of instances in which inflation has breached the Bank's 1-to-3–percent target band. These instances are of particular interest, as "significant" deviations outside the target band mean that the Bank must actively bring inflation back to target.

It may not be appropriate, however, to apply the same 1-to-3–percent threshold to CPIX, as core inflation is less volatile by definition. As evident from Figure 4.2, CPIX breaches the target band much less frequently than All-Items CPI, obscuring any comparative analysis between the policy horizons associated with the two measures. Instead, the policy horizons for core inflation will be evaluated for the months in which All-Items CPI inflation breaches the target band. Under this criteria, the sample of pertinent policy horizons corresponding to "significant" deviations from target is reduced to 70 months for both measures.

Given the specification of inflation as an AR(13) model, the estimated sequences of conditional forecasts will only achieve convergence toward the unconditional mean of the autoregressive process — derived from the data — rather than the nominal 2–percent target. Accordingly, Eqs. (4.11) and (4.13) can be rewritten as follows:

$$|\hat{\pi}_{t+j|t} - \hat{\pi}| \quad < \quad \mu \qquad \forall j \geq N(\hat{\rho}, \hat{\pi}, \pi_t \ldots \pi_{t-13}; \mu) \tag{4.16}$$

$$|\hat{\pi}_{t+j|t} - \hat{\pi}| \quad < \quad \mu \qquad \forall j \geq N(\hat{\rho}, \hat{\pi}, \pi_t \ldots \pi_{t-13}; \mu(\chi, t)) \tag{4.17}$$

where $\hat{\pi}$ is the unconditional mean of the stationary AR(13) process estimated using ordinary least squares, and $\hat{\rho}$ denotes the estimated model parameters. The unconditional means for All-Items CPI and CPIX are 1.98 and 1.92 for the full sample.[9] This feature of the model reveals an additional shortcoming to the recursive approach, as a new unconditional mean must be estimated for each addition to the sample. This gives rise to variation in the tolerance bands within which policy horizons are identified, and this variation can be quite significant over periods of high volatility. In practice, this means that some policy horizons

[9]Note that these are the means of AR(13) processes *estimated* using the inflation data discussed in Section 4.2, and differ from the sample means of the inflation series themselves.

will be associated with unconditional means that are far from the 2–percent inflation target, which is inconsistent with the objective of my analysis.

As discussed in Section 4.1, the choice of the tolerance level for convergence has a significant impact on the length of policy horizons — specifically, a stricter tolerance threshold will produce longer policy horizons, and vice versa. I consider two approaches to convergence — *absolute* convergence associated with the parameter μ, under which forecasted inflation must permanently return to within $\pm\mu$ percentage points of the 2–percent target; and *relative* convergence associated with the parameter χ, under which forecasted inflation must permanently return to within $\pm(1-\chi)$ percent of the deviation from target observed at the beginning of the forecast horizon.[10]

In both cases, the framework employed in this study does not identify an optimal tolerance level. Instead, I propose a small set of values across which I will present policy horizon estimates. In their structural estimation of the optimal policy horizons for Canada, Cayen et al. (2006) and Basant-Roi and Mendes (2007) use an absolute tolerance level of $\mu = 0.10$. However, the policy horizons in this study are estimated using empirical data and exhibit more volatility than the impulse responses from a structural model. Accordingly, I use a slightly less restrictive starting point of 0.25 percentage points, followed by 0.50 and 0.75.[11]

For the relative tolerance criteria χ, I start with an arbitrarily high threshold for convergence at 0.90 (forecasted inflation must fall to within ±10% of the initial deviation), followed by 0.75 and 0.60. To illustrate, consider convergence given $\chi = 0.90$, with All-Items CPI inflation observed at 6.55 percent in 2022M4 and an unconditional mean $\hat{\pi}$ of 1.98 over the full sample. Convergence requires 90% of the initial deviation to have permanently dissipated, which in this example is equal to $(6.55 - 1.98) * 0.10 = 0.46$. Accordingly, inflation is considered to have returned to target if forecasted inflation does not exceed the tolerance band of 1.98 ± 0.46 percentage points for the remainder of the forecast horizon. I view this range of ± 0.46 percentage points as an "effective" tolerance level, denoted μ_ϵ. While its role in estimation is analogous to the absolute tolerance level μ, which is a constant, μ_ϵ depends upon the current level of inflation π_t.

[10]See Davig and Foerster (2021); Batini and Nelson (2001b).

[11]Note that μ has no theoretical upper bound, but values greater than or equal to 1 would lead to the undesired result of convergence outside the Bank's 1-to-3–percent target band.

4.6 Results

4.6.1 Policy Horizons Estimates

I report the policy horizons for All-Items CPI and CPIX inflation from 1989M12 to 2022M4. Using the full-sample approach to estimation, I construct conditional within-sample forecasts over a 60-month horizon for each month starting in 2000M1. As the sample terminates in 2022M4, this produces 269 sets of forecasts. Within each set, I identify the policy horizon as the number of months it takes for forecasted inflation to converge to the tolerance band without breaching it for the remainder of the forecast horizon.

As a concrete example, Figure 4.5 shows forecasted inflation starting from 2022M5, 1 month following the last observation of the full sample. With an absolute tolerance μ around the unconditional mean, the forecasts converge to and remain within the tolerance band in 34 months (approximately 11 quarters) for All-Items CPI inflation (panel (a)) and 21 months (7 quarters) for CPIX inflation (panel (b)) — these are the policy horizon estimates associated with 2022M4. As I will later discuss in greater detail, larger deviations of inflation from target — over 4 percentage points in the example of 2022M4 — imply that forecasted inflation will return to target over a longer horizon.

As stated in Section 4.2, recursive estimation is initiated with *t=1* fixed at 1991M1 and an end date of 1999M12. I then construct conditional forecasts over a 60–month horizon from 2000M1 onward.[12] For each month added to the recursive sample, the forecast horizon is also shifted forward by the same increment, so that the set of forecasts constructed for the last recursive estimate is identical to those for the full sample estimation. Policy horizons are identified in the same manner for each set of out-of-sample forecasts.

Finally, the split sample approach estimates four distinct models of inflation corresponding to the tenures of Governors Dodge, Carney, Poloz, and Macklem. I omit the sample associated with Governor Macklem's tenure since there are not enough observations to estimate a well-behaved model. For the three other models, within-sample forecasts are constructed for each month in the split sample, with policy horizons again identified in the same manner.

Given the disadvantages of recursive estimation as discussed in Sections 4.3 and 4.5, I focus on results obtained from the full sample estimation. For reference, I report summary statistics from the recursive estimation in Appendix A. Subsequently, as an extension of the analysis for the full sample results, I explore split sample policy horizons in Section 4.6.3.

[12]The forecasts are dynamic in that forecasted values, rather than the actual data, are used to produce subsequent forecasts.

Figure 4.5: Estimated Policy Horizon (2022M4)

(a) Total CPI Inflation ($\mu = 0.50$)

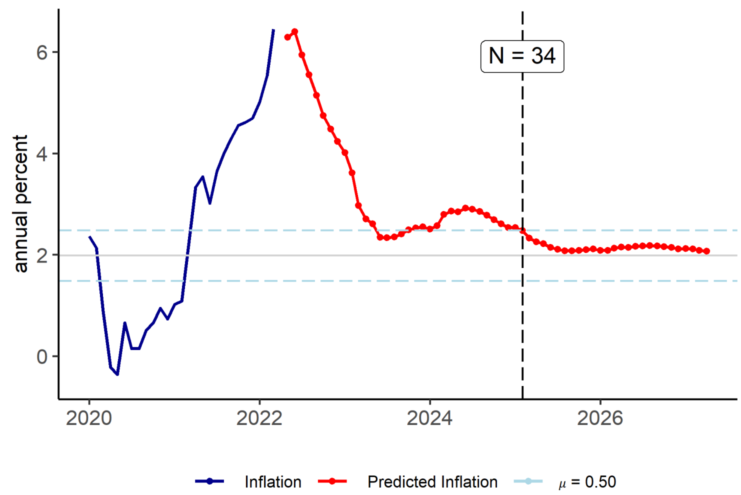

(b) CPIX Inflation ($\mu = 0.25$)

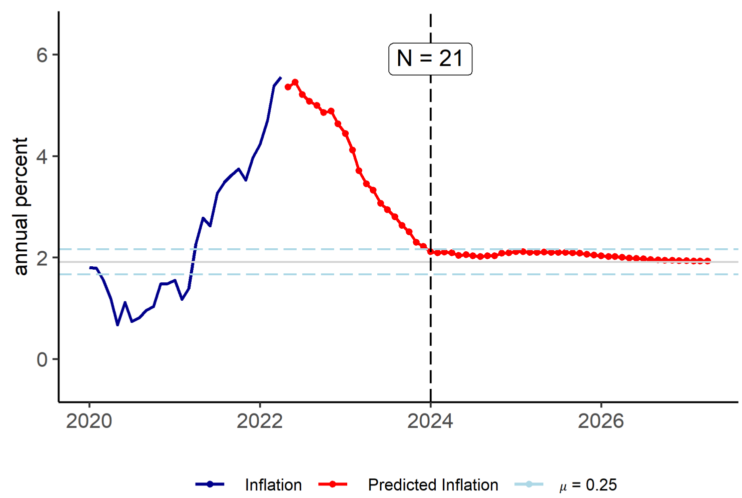

4.6.1.1 All-Items CPI Inflation

Summary statistics for the estimated policy horizons from the full sample estimation are reported in Table 4.2 for All-Items CPI inflation. For each measure of inflation, there are two sets of estimates corresponding to absolute and relative convergence. As discussed in Section 4.5, the summary statistics are only calculated using the subset of policy horizons for which the initial deviation of All-Items CPI inflation from the 2–percent target exceeds 1 percentage point in absolute value.

As expected, a wider tolerance — whether through larger values of μ or smaller values of χ — leads to shorter policy horizons on average. Beginning with absolute convergence and All-Items CPI inflation in panel (a) of Table 4.2, the mean policy horizon is 30 months (10 quarters) when $\mu = 0.25$, with a standard deviation of 5.4 months (approximately 2 quarters). Compared to the Bank of Canada's typical target horizon of 6 to 8 quarters, the estimated mean falls somewhat above this range. A limited range of plus or minus one standard deviation is approximately 25 to 35 months (8 to 12 quarters). Panel (a) of Figure 4.6 illustrates that although the distribution of estimates is clearly asymmetric, this one-standard deviation range covers the majority of observations. With a small cluster of very short horizons, the full range of estimates extends to 9 to 37 months (approximately 3 to 12 quarters).

Table 4.2: Policy Horizons for All-Items CPI Inflation

(a) Absolute Convergence				(b) Relative Convergence			
μ	0.25	0.50	0.75	χ	0.90	0.75	0.60
				$\bar{\mu}^1$	0.17	0.43	0.69
Mean	30.43	14.27	9.04	Mean	34.14	17.26	10.03
Median	32.00	12.00	9.00	Median	34.00	12.50	10.00
Std. Dev.	5.40	8.27	4.57	Std. Dev	3.27	8.64	3.48
Min	9.00	4.00	1.00	Min	28.00	8.00	4.00
Max	37.00	34.00	30.00	Max	49.00	36.00	25.00
Obs.	70	70	70	Obs.	70	70	70

[1] Mean tolerance for convergence ($\hat{\pi} \pm \bar{\mu}_\epsilon$) conditional on χ

Increasing μ from 0.25 to 0.50 shortens the mean policy horizon by more than half to 14 months (approximately 5 quarters), while the standard deviation increases significantly to 8.3 months (approximately 3 quarters). Compared to the Bank of Canada's typical target horizon of 6 to 8 quarters, the estimated mean falls slightly below this range. Combined with the mean policy horizon for $\mu = 0.25$, this implies that the Bank's *de facto* tolerance level — given my particular model of inflation — must lie somewhere between $0.25 < \mu < 0.50$. A limited range of plus or minus one standard deviation is approximately 6 to 22 months

(2 to 7 quarters). As evident from panel (b) of Figure 4.6, however, this range reveals little information about the distribution of estimates. While the modal policy horizon is 12 months and most observations fall between 4 and 13 months, there is a subset of longer estimates distributed around 29 months. This bimodal distribution, and the significant distance between the two modes, lead to much greater variation compared to the estimates from $\mu = 0.25$. The full range of policy horizons is 4 to 34 months (approximately 1 to 11 quarters). Increasing μ further to 0.75 brings the mean policy horizon down to 9 months

Figure 4.6: Estimated Policy Horizons for All-Items CPI Inflation — Absolute Tolerance

(a) $\mu = 0.25$

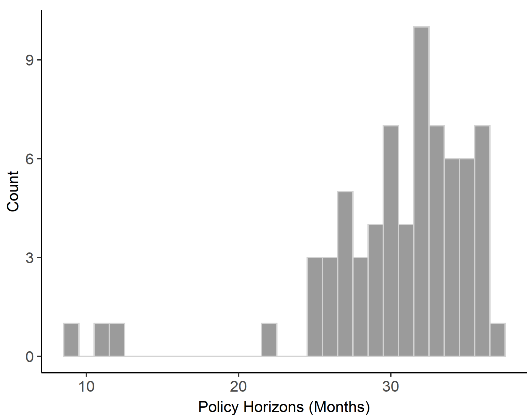

(b) $\mu = 0.50$

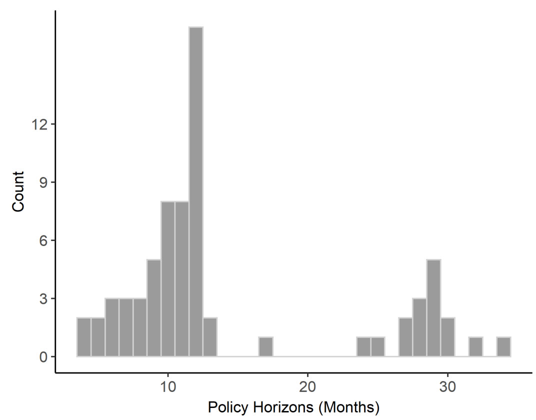

(c) $\mu = 0.75$

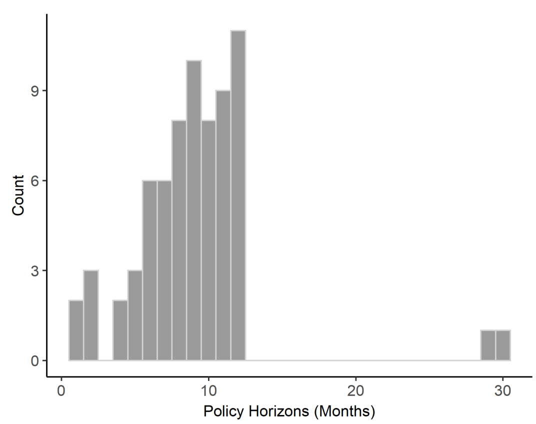

(3 quarters). The standard deviation falls to 4.57 (approximately 2 quarters). A limited range of plus or minus one standard deviation is approximately 4.5 to 13.5 months (1.5 to 4.5 quarters). Panel (c) of Figure 4.6 illustrates that similar to the case of $\mu = 0.25$, this one-standard deviation range covers the majority of observations. All but two policy horizon estimates are 12 months or shorter. Given the outliers at 29 and 30 months, the full range of estimates is 1 to 30 months (approximately 0 to 10 quarters).

In panel (b) of Table 4.2, I present summary statistics for policy horizons estimated under relative convergence for All-Items CPI inflation. The mean of the effective tolerance defined in Section 4.5 — denoted here as $\bar{\mu}_\epsilon$ — is reported as an additional summary statistic. For $\chi = 0.90$, the mean effective tolerance is 0.17. The mean policy horizon is 34 months (approximately 11 quarters), well exceeding the Bank's typical 6-to-8–quarter target range. With a standard deviation of 3.27 (approximately 1 quarter), a limited range of plus or minus one standard deviation is approximately 30 to 37 months (10 to 12 quarters). Panel (a) of Figure 4.7 illustrates that when omitting the small number of outliers beyond 40 months, this limited range of estimated policy horizons captures a reasonably symmetric portion of the distribution. The full range of estimated policy horizons is 28 to 49 months (approximately 9 to 16 quarters).

Figure 4.7: Estimated Policy Horizons for All-Items CPI Inflation — Relative Tolerance

(a) $\chi = 0.90$

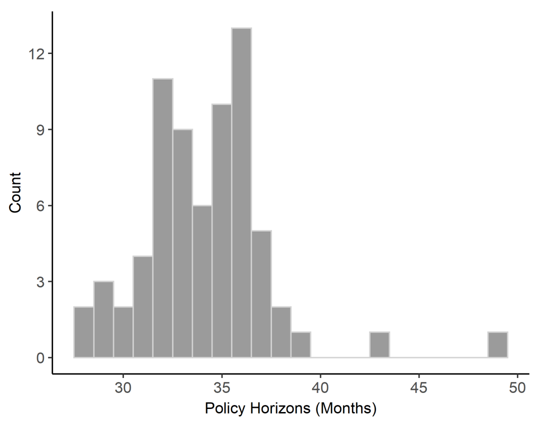

(b) $\chi = 0.75$

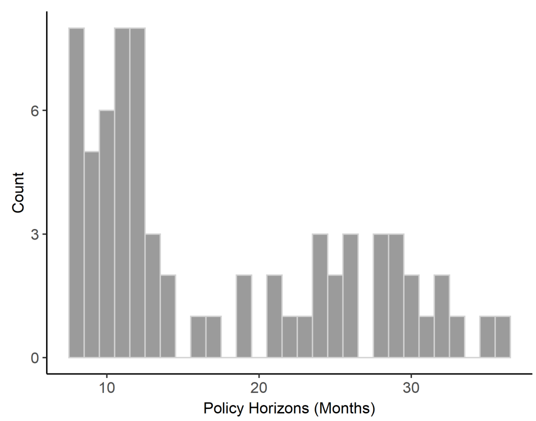

(c) $\chi = 0.60$

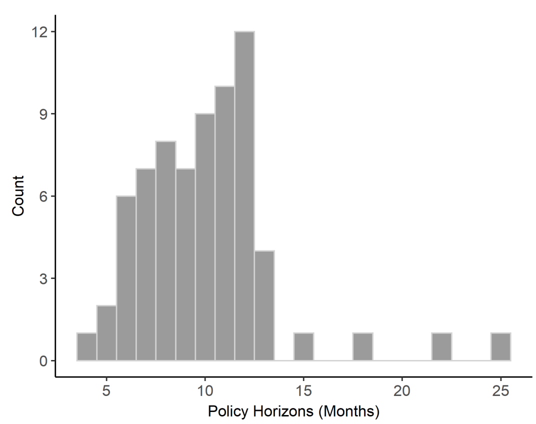

Lowering χ to 0.75 increases the mean effective tolerance to 0.43 percentage points. The mean policy horizon is 17 months (approximately 6 quarters), slightly below the Bank's typical target horizon. Therefore, the relative tolerance corresponding to the 6-to-8–quarter target horizon must fall between $0.75 < \chi < 0.90$, likely near the lower bound of the interval. With a standard deviation of 8.64 months (approximately 3 quarters), a limited range of plus or minus one standard deviation is approximately 8 to 26 months (3 to 9 quarters). Panel (b) of Figure 4.7 illustrates that while this limited range covers the majority of observations, the distribution of estimates is skewed to the right. The full range of estimates is 8 to 36 months (approximately 3 to 12 quarters).

Lowering χ further to 0.60 increases the mean effective tolerance to 0.69 percentage points. The mean policy horizon is 10 months (approximately 3 quarters), well below the Bank's typical target horizon. With a standard deviation of 3.48 (approximately 1 quarter), a limited range of plus or minus one standard deviation is approximately 6.5 to 13.5 months (2 to 4.5 quarters). Again, panel (c) of Figure 4.7 illustrates that this limited range covers the majority of observations, with a small number of outliers beyond 13 months. The full range of estimates is 4 to 25 months (approximately 1 to 8 quarters).

4.6.1.2 CPIX Inflation

I present summary statistics for estimated policy horizons for CPIX inflation in Table 4.3. Under absolute convergence in panel (a), the mean policy horizon is 13 months (approximately 4.5 quarters) when $\mu = 0.25$, with a standard deviation of 5.20 (approximately 2 quarters). A limited range of plus or minus one standard deviation is approximately 8 to 19 months (3 to 6 quarters). The full range of estimates is 1 to 21 months (approximately 0 to 7 quarters), implying that the distribution of estimates is asymmetric.

Table 4.3: Policy Horizons for CPIX Inflation

(a) Absolute Convergence				(b) Relative Convergence			
μ	0.25	0.50	0.75	χ	0.90	0.75	0.60
				$\bar{\mu}^1$	0.08	0.19	0.31
Mean	13.43	8.34	5.33	Mean	29.29	19.20	16.03
Median	15.00	10.00	1.00	Median	23.00	18.00	15.00
Std. Dev.	5.20	6.10	5.46	Std. Dev.	12.05	7.30	6.54
Min	1.00	1.00	1.00	Min	14.00	8.00	4.00
Max	21.00	19.00	17.00	Max	61.00	55.00	49.00
Obs.	70	70	70	Obs.	70	70	70

[1] Mean tolerance for convergence ($\hat{\pi} \pm \bar{\mu}_\epsilon$) conditional on χ

Increasing μ to 0.50 brings the mean policy horizon down to 8 months (approximately 3 quarters), with a standard deviation of 6.10 months (approximately 2 quarters). A limited range of plus or minus one standard deviation is approximately 2 to 14 months (0 to

5 quarters). The full range of estimates is marginally wider at 1 to 19 months (0 to 6 quarters), so the distribution of estimates is skewed to the right.

Further increasing μ to 0.75 brings the mean policy horizon down to just 5 months (approximately 2 quarters), with a standard deviation of 4.57 months (approximately 1.5 quarters). A limited range of plus or minus one standard deviation is 0.5 to 10 months (approximately 0 to 3 quarters). The full range of estimates is 1 to 17 months (approximately 0 to 6 quarters), so the distribution of estimates is again asymmetric and skewed to the right.

Compared to the estimated policy horizons under absolute tolerance for All-Items CPI inflation, those for CPIX inflation are notably shorter for each value of μ. This is expected as core inflation has generally experienced much smaller deviations from target (see Figure 4.2). Additionally, since I estimate policy horizons have for some instances in which All-Items CPI inflation breached the 1-to-3–percent target band but CPIX inflation did not, the forecasts of CPIX inflation will generally converge more quickly.

In panel (b) of Table 4.3, I present summary statistics for the policy horizon estimates of CPIX inflation under relative convergence. Due to the smaller deviations of core inflation from the 2–percent target compared to All-Items CPI inflation, the mean effective tolerances $\bar{\mu_\epsilon}$ are much narrower than those from Table 4.2. This naturally leads to much longer policy horizon estimates compared to those under absolute tolerance, as fluctuations of the non-monotonic inflation forecasts are more likely to breach narrower tolerance bands and extend the policy horizon. For $\chi = 0.90$, the mean effective tolerance is 0.08. The mean policy horizon is 29 months (approximately 10 quarters), with a standard deviation of 12 months (4 quarters). A limited range of plus or minus one standard deviation is 17 to 41 months (approximately 6 to 14). The full range of estimates is 14 to 61 months (approximately 5 to 20 quarters).[13]

For $\chi = 0.75$, the mean effective tolerance is 0.19. The mean policy horizon is 19 months (approximately 6 quarters), with a standard deviation of 7.30 months (approximately 2 quarters). A limited range of plus or minus one standard deviation is 12 to 26.5 months (approximately 4 to 9 quarters). The full range of estimates is 8 to 55 months (approximately 3 to 18 quarters).

For $\chi = 0.60$, the mean effective tolerance is 0.31. The mean policy horizon is 16 months (approximately 5 quarters), with a standard deviation of 6.54 months (approximately 2 quarters). A limited range of plus or minus one standard deviation is 9.5 to 22.5 months (approximately 3 to 7.5 quarters). The full range of estimates is 4 to 49 months (approximately 1 to 16 quarters).

[13]Since the forecast horizon is limited to 60 periods, at least 1 observation within the sequence of forecasts fails to converge in 4 of the 6 cases under relative convergence. In these instances, the model reports a terminal policy horizon length of 61 months but the true horizons may be longer.

While the policy horizons estimated for CPIX inflation provide a useful comparison, the Bank of Canada's inflation-targeting mandate — including the horizon over which it seeks to return inflation to target — is not oriented around core inflation. For this reason, my analysis of flexibility will focus on the policy horizons estimated for All-Items CPI inflation. The discussion will also be limited to absolute tolerance, as this approach more closely aligns with the Bank's stated objective of returning inflation exactly to target regardless of how far inflation has deviated from 2 percent. Comparatively, the relative approach introduces an unnecessary degree of uncertainty to the Bank's communication strategy, which may undermine credibility and potentially weaken inflation expectations. Finally, to refrain from repetition, I will consider only the estimated policy horizons for $\mu = 0.50$.

4.6.2 Evidence of Flexibility

Figure 4.8 plots the percentage deviations of inflation from the 2–percent target against the estimated policy horizons for the 70 months between 2000M1 and 2022M4 in which All-Items CPI inflation exceeded the 1-to-3–percent target band. The 2–percent target is marked as the vertical black line corresponding to a deviation of zero, with the target band marked as dashed blue lines on either side. To better understand these results, it is worth briefly revisiting the interpretation of flexibility pertinent to the empirical analysis.

Section 4.1.2 establishes a very narrow interpretation of flexibility, under which the central bank conducts monetary policy to return inflation to target over a horizon that varies according to the current level of inflation. This interpretation is largely consistent with the concept of broad flexible inflation targeting, defined in Section 2, in which the appropriate policy horizon, $\bar{h}(E_t s_{t+h})$, varies over time based on the central bank's expectation for the path of the state variable s_t over its forecast horizon. The key distinction in the empirical analysis is that the model is not conditioned on variables other than inflation — such as financial stability or full employment — that may influence monetary policy. Given this interpretation, Figure 4.8 shows that the estimated policy horizons are flexible, as there is a clear relationship between current inflation π_t — or more precisely, the distance of π_t from the 2–percent inflation target — and the lengths of policy horizons. As expected, this relationship appears symmetric, as both negative and positive deviations are associated with longer policy horizons as inflation moves further from target.[14]

There is, however, a distinct gap between the clusters of shorter and longer policy horizons. For absolute deviations greater than 1 percentage point, policy horizons either fall between 4 and 13 months, or else above 24 months. The only outlier is an estimate of 17 months

[14]Since dynamics of inflation characterized by the AR(13) model are symmetric, the resulting policy horizons must be also symmetric. However, the deviation-horizon pairs in Figure 4.8 are not perfectly symmetric due to the fact that for a given level of current inflation π_t, each policy horizon estimate may be influenced by different values of lagged inflation.

associated with a deviation of approximately -1. In general, the clusters of longer policy horizons correspond to larger absolute deviations from target, but there is some overlap for deviations around ±2 percentage points in which policy horizons are either very short or very long. The reason for this overlap is that beyond current inflation π_t, the lengths of estimated policy horizons are also influenced by the past lags of inflation, $\pi_{t-1} \ldots \pi_{t-12}$. For a given level of current inflation π_t, the policy horizons associated with different observations may vary significantly depending on the persistence of the sequence of shocks experienced over their recent histories. Specifically, while one-off shocks to inflation generally dissipate within a few months, more persistent shocks may cause inflation to remain outside the 1-to-3–percent target band for an extended period and give rise to longer policy horizons.

Figure 4.8: Policy Horizons vs Inflation Deviation from Target

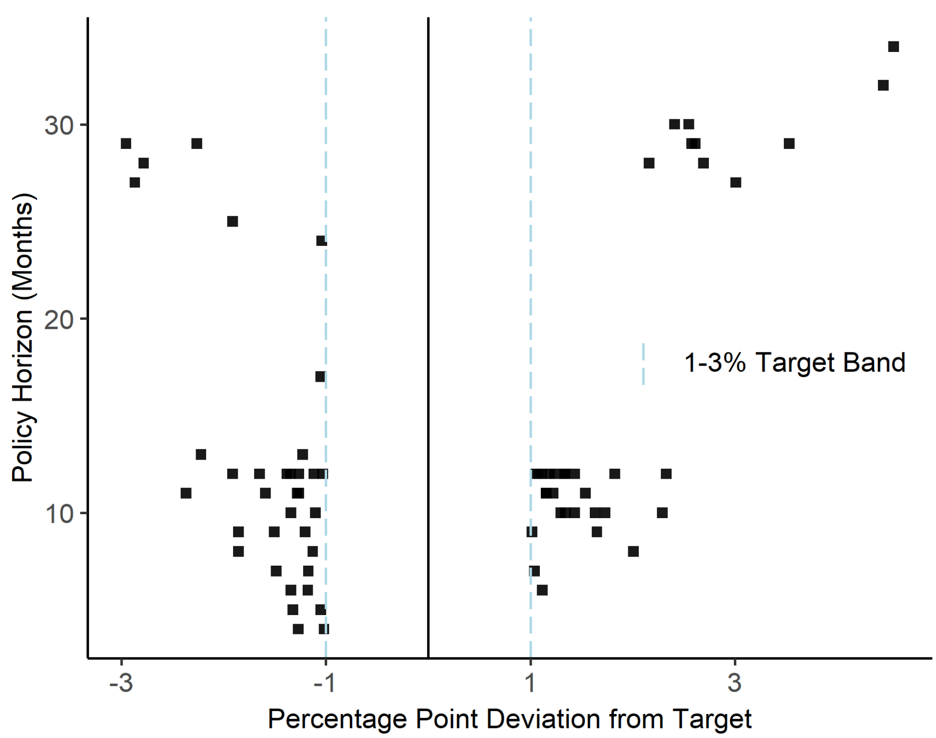

To explore this relationship visually, Figure 4.9 superimposes the policy horizons from Figure 4.8 on the complete time series for All-Items CPI inflation. It is immediately evident that long policy horizons occur in clusters. Specifically, there are three distinct clusters observed around 2003, 2009, and 2022, during which the policy horizon did not fall below 24 months. While each of these may be mapped to specific economic events — a brief but significant shock to commodity prices in early 2003, the financial crisis in 2008–09, and the current bout of high inflation dating back to late 2021 — the empirical framework lacks the sophistication to attribute variation in the lengths of policy horizons to these events. Instead, one may only loosely infer that the history of inflation upon which a particular set of forecasts is conditioned may have been partially determined by the economic circumstances associated with that period. Nevertheless, long policy horizons are only observed when inflation experiences "runs" over which it remains elevated for an extended period. In the three cases identified above, inflation deviated from target by more than 1 percentage point for 6, 8, and 13 consecutive months, respectively.

Figure 4.9: Policy Horizons and All-Item CPI Inflation

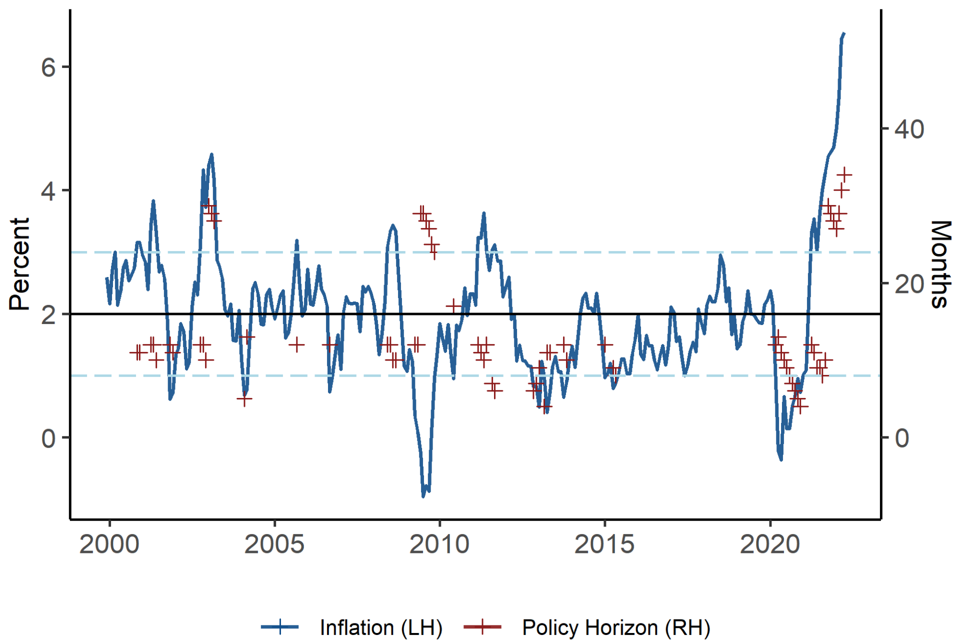

In comparison, for each period in which inflation breaches the target band for just four consecutive months or less, the policy horizon only exceeds 13 months on a single occasion (in 2010M6, with a horizon of 17 months). Policy horizons only extend beyond 2 years — coincidentally, the upper bound of the Bank's typical horizon for monetary policy — when inflation experiences both significant and persistent deviations from target. While it cannot be explicitly identified from the model, this result is consistent with the Bank of Canada's objective of "looking through" transitory shocks to inflation (Bank of Canada, 2016b).

The period associated with the first phase of the Covid-19 pandemic, from 2020M3 to M12, poses an outlier. While inflation remains well below 1 percent for 10 consecutive months, the policy horizon never exceeds 13 months and falls as low as 4 months in 2020M12, when inflation decelerates to just 0.73 percent. As I will reiterate in Section 4.6.3, it is clear that inflation data from the Covid-19 pandemic present significant challenges to the simple model used in this study. Over the full sample, however, there is nevertheless convincing evidence that the narrow interpretation of flexibility applied to the empirical analysis is exhibited by the estimated policy horizons for All-Items CPI inflation. For the periods in which inflation deviates from the 2–percent target by more than 1 percentage point — when monetary policy actions are most relevant — the lengths of the horizons respond in a largely predictable manner to both the current level and recent history of inflation. In the following section, I explore whether the empirical evidence of flexibility is consistent throughout the tenures of the Bank's recent Governors.

4.6.3 Split Sample Policy Horizons

Recall from the simple Fisherian model of inflation in Section 4.1 that the parameter ρ arises from the underlying dynamics of the economy, α from the central bank's monetary policy rule, and σ_ϵ^2 from the variance of the shocks to inflation. For the full sample estimation results discussed in Sections 4.6.1.1 and 4.6.1.2, I assume each of these model parameters to be constant. It may be reasonable, however, to loosen this set of assumptions to accommodate potential changes to both the economic structure and central bank behaviour over time. As discussed in Section 4.3, one approach is to estimate the model using split samples corresponding to the completed tenures of the Bank's recent Governors. Over the full sample period, each Governor experienced unique economic conditions that may have given rise to different dynamics of inflation as well as different shocks. As discussed in Section 3.2, the Bank's monetary policy framework may have also evolved. By dividing the full sample into distinct periods and estimating a unique model for each period, I allow the parameters to vary between each model, potentially giving rise to different distributions of policy horizons.

The split samples used for estimation are 2001M2 to 2008M1 (84 months), when the Bank was led by Governor David Dodge, 2008M2 to 2013M5 (64 months), when it was led by Governor Mark Carney, and 2013M6 to 2020M2 (81 months), when it was led by Governor Stephen Poloz. While Governor Poloz did not leave the Bank until 2020M6, the inclusion of the first months of the Covid-19 pandemic causes the time series to become non-stationary, leading to unstable inflation forecasts and strange policy horizon estimates. More generally, the period from 2020M3 onward marks a significant structural break for inflation and the Canadian economy as a whole, the full impact of which is yet to materialize. While it is likely that the Bank's behaviour has changed in response to the new dynamics of inflation, it is difficult to accurately estimate those dynamics before the shock stabilizes into a new steady state. For this reason, I omit 2020M3 to M6 from the end of the sample corresponding to Governor Poloz's tenure, and omit the sample corresponding to current Governor Tiff

Macklem's tenure altogether.

Table 4.4 presents summary statistics for the policy horizons estimated from split samples, using All-Items CPI inflation and $\mu = 0.50$. Consistent with Section 4.6.1, I consider only the observations for which All-Items CPI inflation deviates from the 2–percent target by more than 1 percentage point.

Table 4.4: Split Sample Policy Horizons

	Dodge (2001M2–2008M1)	Carney (2008M2–2013M5)	Poloz (2013M6–2020M2)
Mean	24.87	11.16	4.40
Median	25.00	13.00	5.00
Std. Dev.	7.80	5.62	1.82
Min	13.00	1.00	2.00
Max	38.00	19.00	6.00
Obs.	15	25	5

Full-sample AR(13) model, All-Items CPI inflation, $\mu = 0.50$.

The policy horizons estimated over Governor Dodge's tenure are the longest out of the three split samples, with a mean horizon of 25 months (approximately 8 quarters). This falls marginally above the Bank's typical 6-to-8–quarter horizon. The full range of horizons, however, is much wider at 13 to 38 months (approximately 4 to 13 quarters). Over Governor Carney's tenure, the mean policy horizon falls by more than half to 11 months (approximately 3 quarters) — slightly below the Bank's typical 6-to-8–quarter horizon. The full range of horizons is 1 to 19 months (approximately 0 to 6 quarters), indicating that the distribution of estimated horizons has shifted significantly to the left of the distribution associated with Governor Dodge's tenure. Finally, over Governor Poloz's tenure, the mean policy horizon is 4 months (approximately 1 quarter), with a range of just 2 to 6 months (approximately 0.5 to 2 quarters). The two latter sets of estimates imply that the absolute tolerance of 0.50 is higher than the *de facto* tolerance that corresponds to the Bank's typical 6-to-8–quarter horizon. As established in Section 4.6.1, this tolerance lies somewhere between 0.25 and 0.50, so the policy horizons presented in Table 4.4 have likely been underestimated. Nevertheless, the split sample estimates provide a useful comparison against the full sample policy horizons, while also revealing significant differences in the distributions of policy horizons associated with each Governor's tenure.

Figures 4.10, 4.11, and 4.12 present the policy horizons estimated using the split samples corresponding to the tenures of Governors Dodge, Carney, and Poloz, respectively. I superimpose the full sample policy horizons over the split sample estimates. Given that the split sample and full sample inflation forecasts are conditioned upon the same recent history of inflation, one could attribute any discrepancy between the policy horizon for a given month to differences between the estimated inflation processes. Recall from Section 4.3 that the parameters estimated from the full sample reflect the average inflation process from 1989M1 to 2022M4, and those estimated from each split sample reflect the average inflation process for that particular period. While the short-term dynamics observed over the split samples are incorporated into the full sample as part of the long-term average, there may nevertheless be notable differences between the resulting policy horizons. Specifically, for each pair of policy horizon estimates associated with the same observation of inflation π_t, the vertical distance may be attributed to differences in the dynamics of inflation, characterized by the estimates of ρ, the variance of the shocks to inflation σ_ϵ^2, or changes in central bank behaviour implied by α. However, given the limitations of my empirical framework, it is impossible to disentangle the specific source of variation between the split sample and full sample distributions of policy horizons. Similarly, the same qualification applies to the differences between the distributions of policy horizons across split samples. These differences may be attributed to the aggregate impact of potential changes in ρ, α, σ_ϵ^2 — and any endogenous responses between the parameters — but in most cases, it is impossible to determine the precise impact of each determinant.

Figure 4.10: Policy Horizons vs Deviation from Target (Dodge Sample)

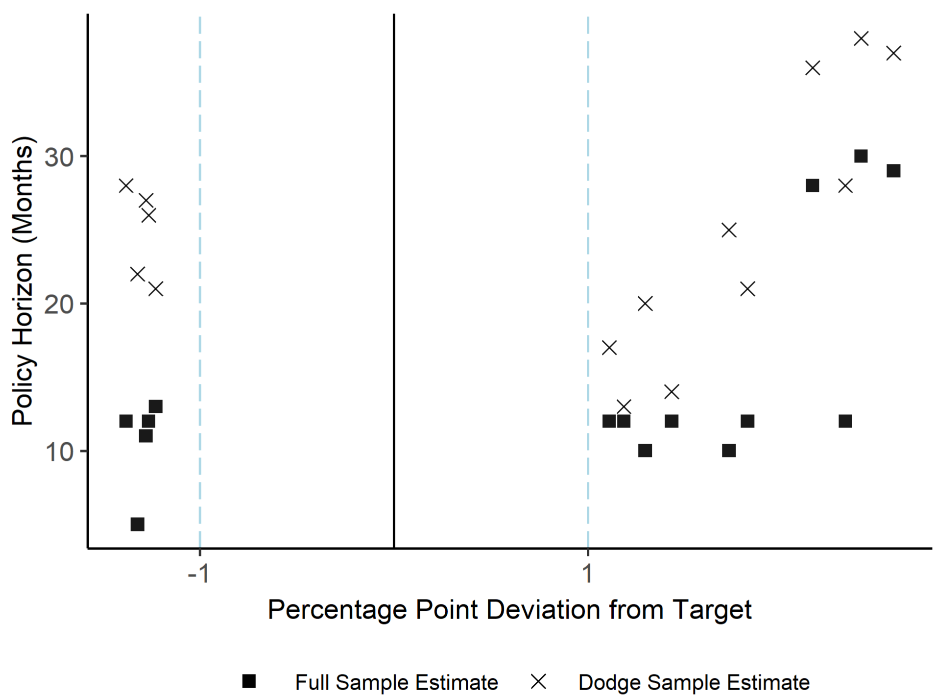

Figure 4.10 clearly shows that over Governor Dodge's tenure, policy horizons estimated using the split sample are considerably longer than those from the full sample. In other words, conditional on the inflation process estimated over 2001M2 to 2008M1, forecasted inflation returns to target more slowly than they do under the inflation process estimated over the full sample. On average, the split sample policy horizons are 10 months longer but follow a pattern very similar to their full sample counterparts. Despite the difference in lengths, the type of flexibility identified in Figure 4.8 is clearly retained in the split sample estimates. The cluster of long policy horizons in the top right corner can be traced to a brief bout of high inflation from late 2002 to early 2003, corresponding to spikes in gasoline, heating oil, and natural gas prices amidst tightening global energy markets (Bank of Canada, 2003). One key difference, however, is that while the split sample policy horizons consistently increase in length as inflation moves further above target, the full sample horizons abruptly skip from 12 to 28 months (see Section 4.6.2). While this change may be driven by differences in the estimates of ρ between the split sample and full samples models, or a particular value of α associated with Governor Dodge's tenure, identification of the precise mechanism is beyond the scope of my empirical framework.

Figure 4.11: Policy Horizons vs Deviation from Target (Carney Sample)

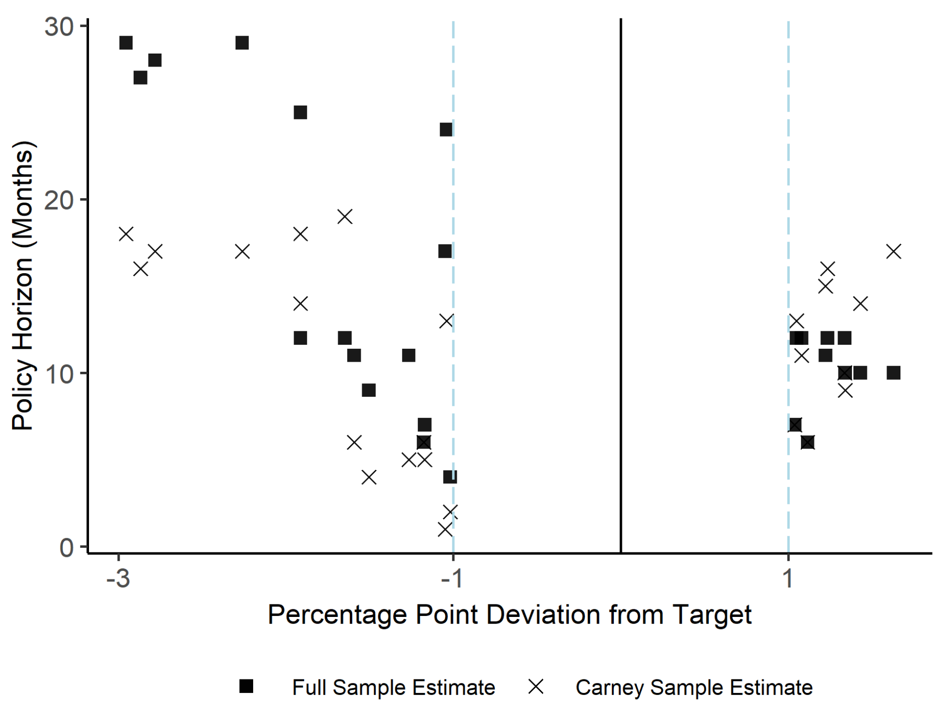

Figure 4.11 shows that over Governor Carney's tenure, the policy horizons estimated using

the split sample are 5.6 months shorter on average than those from the full sample. Again, evidence of flexibility is retained in the split sample estimates, as the policy horizons generally increase in length as inflation moves further below target. The longest policy horizons, which are located in the top left corner, are associated with a period of low inflation (and briefly deflation) in mid-to-late 2009 at the tail end of the global financial crisis.

Figure 4.12: Policy Horizons vs Deviation from Target (Poloz Sample)

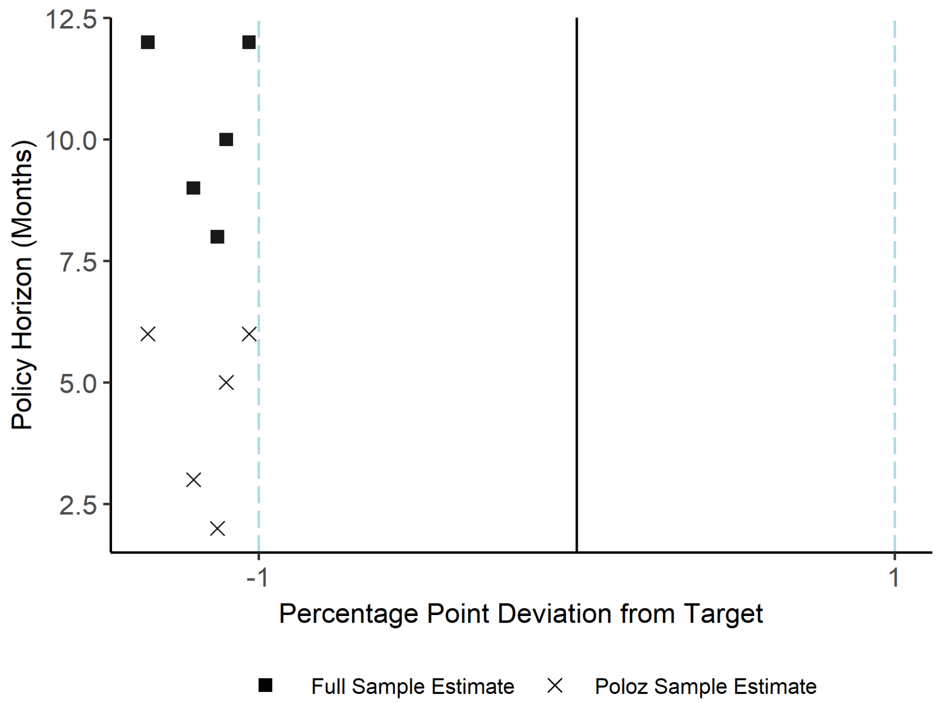

A somewhat surprising result is that split sample horizons associated with Governor Dodge's tenure are significantly longer on average than those associated with Governor Carney's tenure. As discussed in Section 3.2, the 2006 *Renewal* that was in effect throughout Governor Dodge's tenure reiterates that the typical 6-to-8–quarter horizon over which the Bank seeks to return inflation to target remained generally appropriate. In contrast, the subsequent 2011 *Renewal*, published during Governor Carney's tenure, places much greater emphasis on the use of longer policy horizons. Since inflation falls outside the 1-to-3–percent target band much more frequently over Governor Carney's tenure (25) than Governor Dodge's (15), and the average deviation in the former sample is larger in absolute value, this implies that the variance of the shocks to inflation — σ_ϵ^2 — is likely higher for the model estimated over Governor Carney's tenure. As discussed in Section 4.1.2, an increase in σ_ϵ^2 is expected to lead to longer policy horizons on average. The inconsistency between the theoretical and empirical results implies that the other parameters, ρ and α, must be considerably different between the two split samples. Again, however, the precise impact of potential changes to each of the parameters and their interactions cannot be identified.

Finally, Figure 4.12 shows that over Governor Poloz's tenure, the policy horizons estimated using the split sample are, on average, 5.8 months shorter than those from the full sample. This set of estimates is notably shorter than the policy horizons associated with both Governors Dodge and Carney. These differences may be attributed to the lower volatility in inflation experienced over Governor Poloz's tenure relative to the full sample, as well as the tenures of the two previous Governors. With the Covid-19 observations removed, the shocks to inflation experienced over this period are very minor, as inflation only breached the 1-to-3–percent target band five times compared to 15 under Governor Dodge and 25 under Governor Carney. Recall from Section 4.1.2 the variance of inflation σ_π^2 depends upon each of the model parameters, ρ, α, and σ_ϵ^2. Therefore, while inflation was more stable under Governor Poloz compared to his two predecessors, leading to shorter policy horizons, this could have been driven by a particular approach to monetary policy through α, a unique set of dynamics for inflation through ρ, a less volatile sequence of shocks to inflation through σ_ϵ^2, or any combination of these factors and their endogenous responses to each other.

In general, however, the economic experience over Governor Poloz's tenure did not warrant significant monetary policy actions, making it difficult to assess whether the estimated policy horizons exhibit flexibility. The narrow type of flexibility — that the lengths of policy horizons vary based on the distance of current inflation from target — is clearly evident in the split sample estimates over the tenures of both Governors Dodge and Carney. The same cannot be said for the estimates from Governor Poloz's tenure. Taken together, it is unclear from the empirical evidence whether the Bank's approach to flexibility has materially changed over its inflation targeting period, at least up to the start of the Covid-19 pandemic in early 2020. The full impact of the Covid-19 experience — including the ongoing episode of high inflation — on inflation expectations, central bank credibility, and flexibility falls well beyond the scope of this study but marks an important avenue for future research.

4.6.4 Robustness Checks

4.6.4.1 Residual Autocorrelation

One concern discussed in Sections 4.2 and 4.4 is the high degree of persistence inherent to the inflation data. While I err toward a parsimonious model in the study, it is not surprising that the residuals from the model for both All-Items CPI and CPIX were not free of autocorrelation (see Figure 4.4). To evaluate whether the residual autocorrelation impacted the estimation of policy horizons, I re-estimate the full sample model for both All-Items CPI and CPIX inflation using the optimal lags selected from likelihood ratio tests. As discussed in Section 4.4, the lag lengths were 49 for All-Items CPI and 59 for CPIX.

Panels (a) and (b) of Figure 4.13 present the partial autocorrelation functions (PACF) for the full sample residuals for All-Items CPI inflation and CPIX inflation. Compared to the AR(13) residuals in panel (a) of Figure 4.4, the AR(49) model removes most of the autocorrelation and the residuals appear much better behaved. Similarly, the PACF plot for the extended AR(59) model of CPIX inflation shows no evidence of autocorrelation save for a spike close to the 5% confidence level at $p = 50$. While neither of the extended models generates fully i.i.d. residuals, they mark a significant improvement over the residuals from the parsimonious models.

Table 4.5: Select Policy Horizons for All-Items CPI Inflation

Absolute Convergence — Full-Sample Model						
	Base – AR(13)			Extended – AR(49)		
μ	0.25	0.50	0.75	0.25	0.50	0.75
Mean	30.43	14.27	9.06	44.51	15.11	8.66
Median	32.00	12.00	9.00	51.00	12.00	10.00
Std. Dev.	5.40	8.27	4.57	14.89	8.48	3.35
Min	9.00	4.00	1.00	11.00	4.00	1.00
Max	37.00	34.00	30.00	61.00	34.00	12.00
Obs.	70	70	70	70	70	70

Using the extended AR(49) and AR(59) models, I estimate the policy horizons for All-Items CPI and CPIX inflation. I present the corresponding summary statistics in Table 4.5 and Table 4.6 alongside the summary statistics from the AR(13) model for ease of comparison. For the sake of brevity, I report summary statistics for the full sample model under absolute convergence only. For All-Items CPI, the policy horizons estimated from the extended model under $\mu = 0.25$ are noticeably longer than those from the base model. The mean policy horizon is 44 months (15 quarters) compared to 30 months (10 quarters) in the base model, and the standard deviation is nearly three times as large at 14.89. The range of estimates spans 11 to 61 months (approximately 4 to 20 quarters). Recall from Section 4.6.1 that policy horizons terminating at 61 months indicate that the forecast horizon of 60 months is not sufficient for convergence, leading the mean policy horizon to be slightly underestimated. This result is expected, as the longer dynamics in the extended model will accentuate the

Figure 4.13: Partial Autocorrelation Functions of Model Residuals

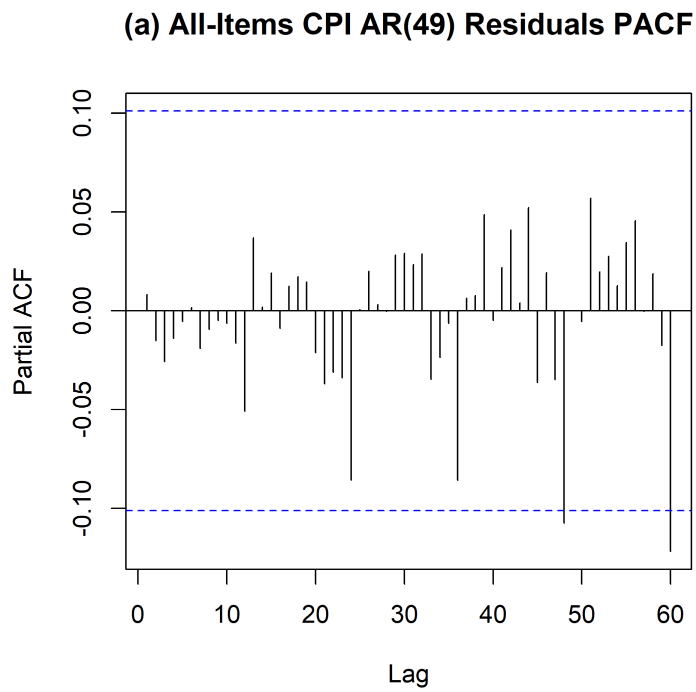

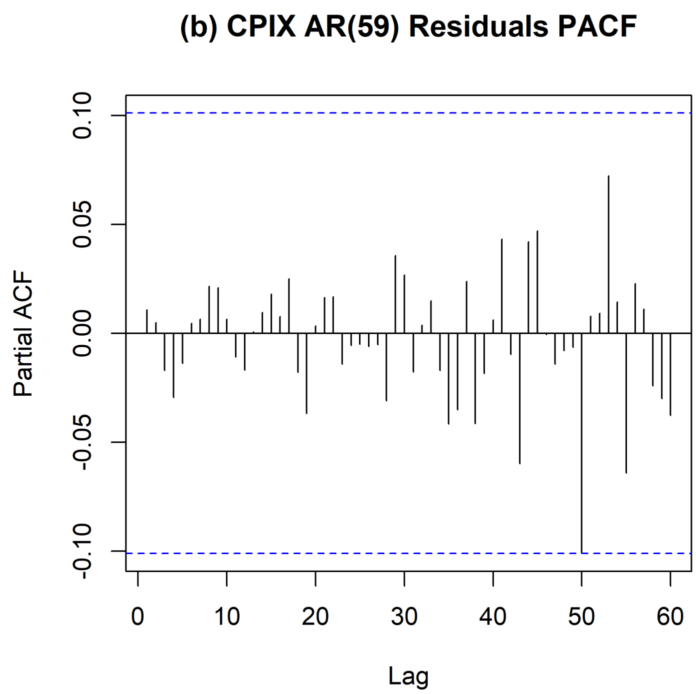

non-monotonicity of forecasts, leading to more frequent breaches of a tight tolerance band and therefore longer policy horizons.

For $\mu = 0.50$, however, the estimated policy horizons between the base and extended models are nearly identical. The mean policy horizon is 15 months (5 quarters), compared to 14 months (approximately 5 quarters) in the base model, and the standard deviation is 8.48 compared to 8.27. In both cases, the range of estimates is 4 to 34 months (approximately 1 to 11 quarters). This result indicates that for the specific set of estimates on which the analysis in Sections 4.6.2 and 4.6.3 are focused — namely, absolute convergence for All-Items CPI inflation with $\mu = 0.50$ — the presence of autocorrelation in the residuals of the AR(13) model do not meaningfully impact the key results.

For $\mu = 0.75$, the estimates from the base and extended models are also similar. The mean policy horizon is 8.66 months compared to 9 months (approximately 3 quarters), and the standard deviation is slightly lower at 3.35. While the range of estimates appears much narrower at 1 to 12 months (0 to 4 quarters) compared to 1 to 30 months (0 to 10 quarters), closer inspection of the AR(13) estimates shows that other than two outliers for which the policy horizons are 29 and 30 months, the remainder of the estimates do not exceed 12 months.

Table 4.6: Select Policy Horizons for CPIX Inflation

Absolute Convergence — Full-Sample Model						
	AR(13)			AR(59)		
μ	0.25	0.50	0.75	0.25	0.50	0.75
Mean	13.43	8.34	5.33	47.94	18.44	12.47
Median	15.00	10.00	1.00	53.00	12.00	3.50
Std. Dev.	5.20	6.10	5.46	12.22	19.24	16.68
Min	1.00	1.00	1.00	2.00	1.00	1.00
Max	21.00	19.00	17.00	61.00	57.00	54.00
Obs.	70	70	70	70	70	70

The policy horizons estimated for CPIX using the extended model bear very little resemblance to those from the base model. The mean policy horizon is more than double the length of the base model estimates for each value of μ, along with significantly larger standard deviations. The notably wider ranges of estimates indicate that fluctuations of forecasts are amplified by the longer dynamics of the extended model, leading to outliers of long policy horizons even for high tolerance levels. However, without a reference point for CPIX analogous to the Bank's target horizon of 6 to 8 quarters for All-Items CPI inflation, there is insufficient information to determine which specification more accurately captures the dynamics of core inflation.

4.6.4.2 Selection of μ

As an additional robustness check, I evaluate whether the relationship between the mean policy horizon and the tolerance level is relatively smooth over the range of possible values for μ. Since the analysis in Section 4.6 is confined to an arbitrary sample of three values of $\mu = 0.25$, 0.50, and 0.75, I have necessarily assumed that the behaviour of the estimated policy horizons will not change idiosyncratically with the choice of μ. In other words, the validity of the choice of μ — particularly the focus on $\mu = 0.50$ in Sections 4.6.2 and 4.6.3 — relies on the assumption that the estimate policy horizons do not behave differently across some threshold value.

To evaluate this assumption, I re-estimate the sequences of in-sample forecasts for the full-sample model of All-Items CPI inflation and CPIX inflation, iterating over all values of μ from 0.01 to 1.00 in increments of 0.01. Recall that for a given value of μ, the AR(13) model is estimated using the full sample period from 1989M12 to 2022M4, then in-sample conditional forecasts are estimated up to 60 months ahead for every observation between 2000M1 and 2022M4 in which All-Items CPI inflation breaches the 1-to-3–percent target band. This generates a sample of 70 policy horizon estimates. In this exercise, I collect the average length of the 70 estimated policy horizons for each iterative value of μ, forming a sample of 100 averages. Figure 4.14 plots the mean policy horizon against the value of μ used for each particular set of estimates. *Ex-ante*, I expect a generally negative relationship between the value of μ and the mean policy horizon, since convergence should occur more quickly for higher tolerance levels.

The relationship between the mean policy horizon and μ closely resembles an exponential decay function for both All-Items CPI and CPIX. For lower values of μ, marginal increases to the tolerance level are associated with relatively large decreases in the mean policy horizon, while for larger values of μ, the mean policy horizon decreases at an increasingly slower rate in response to marginal increases in μ. Framing the relationship in terms of elasticities, the point of "unit" elasticity appears to occur around $\mu = 0.30$ for All-Items CPI and around $\mu = 0.175$ for CPIX. The rationale for the heightened sensitivity of policy horizons to changes around low values of μ is straightforward — given non-monotonic forecasts from the AR(13) model, a tight tolerance band gives rise to long policy horizons in which the forecast sequence appears to converge much earlier but breaches the band towards the end of the horizon. Each incremental increase to the size of the tolerance band reduces the likelihood of such an outlier occurring, which in turn has a potentially significant impact on the mean policy horizon. For higher values of μ, the likelihood of forecasts breaching a wider tolerance band near the end of the forecast horizon is relatively low, as the fluctuations dampen over time. Therefore, incremental increases to the tolerance band will not eliminate long policy horizons to the same extent. At the other end of the forecast horizon, incremental

Figure 4.14: Average Policy Horizon for $\mu \in (0, 1]$

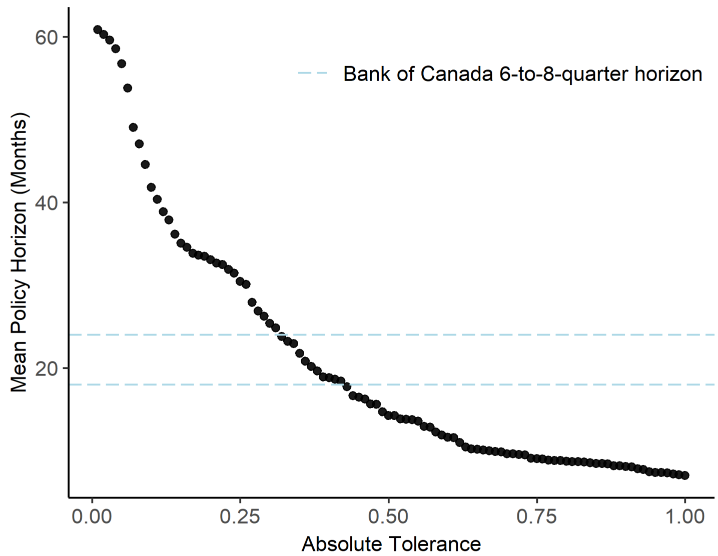

(a) All-Items CPI Inflation

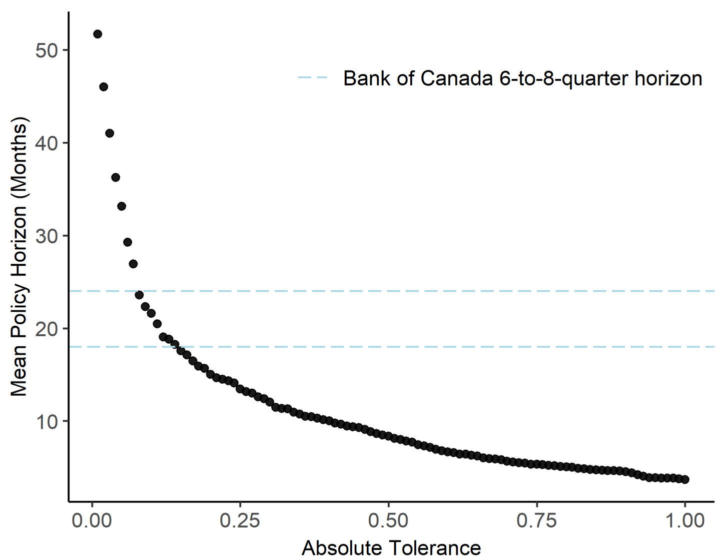

(b) CPIX Inflation

increases to an already wide tolerance band will not necessarily lead to considerably shorter policy horizons, as fluctuations at the beginning of the forecast sequence are much more amplified. This means that larger increases to μ are needed to bring a particular sequence of fluctuations into the tolerance band, leading to the mean policy horizon falling more slowly over incremental changes.

An additional application of this exercise is the identification of the *de facto* tolerance level that matches the Bank of Canada's 6-to-8–quarter target horizon. From Section 4.6.1, the range of potential values was found to be $0.25 < \mu < 0.50$. Using the underlying data from Figure 4.14, the exact value for which the full-sample AR(13) model would produce a mean policy horizon that matches the midpoint of the target range at 7 quarters. For increments of 0.01, the closest candidate value is $\mu = 0.36$, which is associated with a mean policy horizon of 20.83 months.

One potential point of concern can be identified in the plot for All-Items CPI, in which there is a slight "bump" around $\mu = 0.25$. For values of approximately $0.20 < \mu < 0.25$, the mean policy horizon appears to be considerably less sensitive to changes in the tolerance level. However, this cluster of observations is more appropriately viewed as a minor discrepancy than a discontinuity, as the overall trend of the relationship is rather distinct and the magnitude of the deviation from the trend is quite small. Additionally, this "bump" occurs well to the left of the focal value of $\mu = 0.50$, within whose proximity the relationship is adequately stable. It also occurs to the left of $\mu = 0.36$, so there are no discernable discontinuities between the *de facto* tolerance level associated with the Bank's target horizon and the focal value of $\mu = 0.50$ for All-Items CPI. Accordingly, the selection of 0.25, 0.50, and 0.75 appears innocuous.

Chapter 5

Discussion and Conclusion

While the narrative and empirical components of this study differ considerably in scope and method, both contribute to building an understanding of the theory and practice of flexible inflation targeting. To establish the broader context, the narrative analysis first presents a tailored overview of the inflation targeting frameworks followed by modern central banks, illuminating common themes and identifying key differences. In brief, most inflation-targeting central banks in the OECD allow for some degree of flexibility in the horizon over which they seek to return inflation to target. This flexibility is generally exhibited through the lengths of policy horizons, which adjust over time in response to particular economic conditions and shocks to inflation. There is, however, considerable variation in how monetary policy is communicated to the public, reflecting the varying degrees of emphasis on transparency across central banks.

The detailed review of the Bank of Canada then explores the extent to which its inflation targeting framework has evolved since its inception in 1991. Initially, the framework's objective was limited to returning inflation to 2 percent within a target horizon of 6 to 8 quarters. However, given the Bank's experiences through the financial crisis and the Covid-19 pandemic, relevant considerations for monetary policy gradually expanded beyond inflation to include financial stability, maximum sustainable employment, and output stabilization. In turn, the Bank has become increasingly more receptive to longer policy horizons, and to using the flexibility at its disposal to pursue these complementary objectives when conditions allow. How the Bank communicates its policy horizons through the *Monetary Policy Report* has also seen considerable change, reflecting the preferences of different Governors and the economic challenges they confront. As a result, the evolution of transparency in the Bank's communications has been somewhat inconsistent. By and large, however, the Bank's inflation targeting framework has been successful at maintaining low and stable inflation while keeping inflation expectations anchored around the 2–percent target.

The empirical analysis puts forth up-to-date estimates of *de facto* policy horizons for Canada. Focusing on All-Items CPI inflation and an absolute tolerance level of $\mu = 0.50$, I find that when inflation breaches the 1-to-3–percent target band, the associated policy horizons span from 4 to 34 months (approximately 1 to 11 quarters), with a mean of 14 months (approximately 5 quarters). While differing in methodology, this range of estimates is largely consistent with those reported by Cayen et al. (2006) and Basant-Roi and Mendes (2007) — 4 to 11 quarters and 2 to 9 quarters, respectively — which Coletti et al. (2006) cite in support of the Bank's 6-to-8–quarter target horizon. For both the full sample and split sample models, the estimated policy horizons exhibit flexibility in a very narrow manner — the lengths of policy horizons depend upon the distance of current and recent inflation from target.

Additionally, the split sample estimation reveals significant differences between the distributions of policy horizon estimates associated with the tenures of Governors Dodge, Carney, and Poloz. This is expected, as each Governor experienced different economic conditions and sequences of shocks to inflation, which in turn necessitate different policy responses. However, given the limitations of the empirical framework, I cannot determine the extent to which the variation between distributions can be attributed to explicit changes in central bank behaviour rather than structural changes in the economy or different sequences of shocks. Therefore, it is also unclear whether the Bank's evolution toward greater flexibility, as demonstrated through its written framework, can be mapped to distinguishable changes in the distributions of policy horizon estimates over time.

On the broader concept of policy horizons, I find that unusually high or low inflation does not necessarily lead to long horizons. Instead, shorter policy horizons consistently arise from significant but brief deviations of inflation from target. This result may reflect the Bank's commitment to "looking through" transitory shocks to inflation, as established in its monetary policy framework, although this connection cannot be explicitly drawn from my empirical analysis (Bank of Canada, 2016b). Meanwhile, longer policy horizons arise when inflation is not only far from target but has remained far from target for an extended period. In other words, the persistence of policy horizons and inflation are highly correlated, as monetary policy decisions and the expected path of inflation endogenously affect each other.

One of the key challenges to the empirical analysis is the extraordinary degree of volatility that inflation has experienced since early 2020, which amplifies the variance of the time series over the full sample and potentially skews the distribution of policy horizon estimates. While the split sample approach circumvents this issue by terminating the sample at 2020M2, it does so at the expense of estimating policy horizons for Governor Macklem's tenure. It will be worthwhile, therefore, to revisit both the full sample estimation and the split sample model for Governor Macklem once the current episode of high inflation has stabilized.

An additional challenge is the impact of the effective lower bound (ELB) on policy horizons. As monetary policy is constrained by the ELB when both inflation and the policy rate are close to zero, a negative inflation shock will necessarily lead to longer policy horizons compared to an equivalent positive shock when inflation is above target. While the empirical framework implicitly incorporates the presence of the ELB through the inflation data, the model itself is not constrained by this underlying asymmetry. Instead, the conditional forecasts constructed from the model are symmetric, so that negative and positive deviations of inflation from target may give rise to similar policy horizons. As a result, for certain periods when the Bank has been confronted with deflation or very low inflation, the associated policy horizons may have been underestimated.

More broadly, the simplicity of the univariate model limits the empirical analysis. As emphasized in Section 4.6.3, changes in the distribution of policy horizon estimates over time cannot be attributed to specific factors beyond the overall variation in inflation. One reason for this limitation is that additional variables that may determine inflation — such as output growth, exchange rates, and household balance sheets, to name a few — are omitted from the model. By extending the analysis to a vector autoregressive (VAR) framework, these variables can be used to specify a much richer model of inflation, and allow policy horizons to be identified using impulse responses from specific sources of shocks. Alternatively, a general equilibrium approach could leverage structural equations and policy rules to model more complex dynamics, capture relevant asymmetries, and potentially produce more robust policy horizon estimates.

Taken together, each of these limitations offers opportunities to expand upon the empirical framework and provide deeper insights into empirical policy horizons. Toward this objective, Mu and Voss (forthcoming) further explores the themes discussed in this study and pursues more rigorous estimates of policy horizons for Canada.

Bibliography

Andersson, F. N. G., and L. Jonung (2017) *How Tolerant Should Inflation-Targeting Central Banks Be? Selecting the Proper Tolerance Band - Lessons from Sweden*, 2017-2, Working Papers, Lund University, Department of Economics

Bailey, A. (2022) "Letter from the Governor to the Chancellor,"

Banco Central de Chile (2022) "Monetary Policy Report,"

Bank of Canada (1991) "Targets for reducing inflation,"

——— (1993) "Statement of the Government of Canada and the Bank of Canada on Monetary Policy Objectives,"

——— (2001) "Renewal of the Inflation-Control Target,"

——— (2002) "Monetary Policy Report,"

——— (2003) "Monetary Policy Report,"

——— (2006) "Renewal of the Inflation-Control Target,"

——— (2008a) "Monetary Policy Report,"

——— (2008b) "Monetary Policy Report,"

——— (2009) "Monetary Policy Report,"

——— (2011) "Renewal of the Inflation-Control Target,"

——— (2012) "Monetary Policy Report,"

——— (2014a) "Monetary Policy Report,"

——— (2014b) "Monetary Policy Report,"

——— (2016a) "Monetary Policy Report,"

——— (2016b) "Renewal of the Inflation-Control Target,"

——— (2020) "Monetary Policy Report,"

——— (2021a) "Bank of Canada will hold current level of policy rate until inflation objective is sustainably achieved, continues quantitative easing,"

——— (2021b) "Monetary Policy Report,"

——— (2021c) "Monetary Policy Report,"

——— (2021d) "Monetary Policy Report,"

——— (2021e) "Renewal of the Inflation-Control Target,"

——— (2022a) "Bank of Canada increases policy interest rate by 100 basis points, continues quantitative tightening,"

——— (2022b) "Monetary Policy Report,"

——— (2022c) "Monetary Policy Report,"

Bank of England (2013) "Review of the monetary policy framework,"

——— (2022) "Monetary Policy Report,"

Bank of Israel (2014) "Conducting Monetary Policy,"

Basant-Roi, M., and R. R. Mendes (2007) *Should Central Banks Adjust Their Target Horizons in Response to House-Price Bubbles?*, Staff Discussion Paper, Bank of Canada

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

Board of Governors of the Federal Reserve System (2020) "Guide to changes in the 2020 Statement on Longer-Run Goals and Monetary Policy Strategy,"

——— (2022a) "Federal Reserve issues FOMC statement,"

——— (2022b) "Monetary Policy Report,"

Cayen, J.-P., A. Corbett, and P. Perrier (2006) "An Optimized Monetary Policy Rule for ToTEM," *Bank of Canada Working Paper 2006-41*

Central Bank of Iceland (2001) "Declaration on the inflation target and a change in the exchange rate policy,"

Chung, H. T., B. M. Doyle, J. Hebden, and M. Siemer (2020) *Considerations Regarding Inflation Ranges*, 2020-75, FEDS Working Paper

Coletti, D., J. Soledy, and C. Wilkins (2006) "Another Look at the Inflation-Target Horizon," *Bank of Canada Review* pp. 31–37

Davig, T. A., and A. Foerster (2021) "Communicating Monetary Policy Rules," (2021-12)

European Central Bank (2021) "Monetary Policy Strategy,"

——— (2022) "Economic Bulletin,"

Goodfriend, M. (2003) *Inflation Targeting in the United States?*, 9981, NBER Working Paper

Khan, M., L. Morel, and P. Sabourin (2015) *A Comprehensive Evaluation of Measures of Core Inflation for Canada*, Staff Discussion Paper, Bank of Canada

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

Lane, P. R. (2022) “The monetary policy strategy of the ECB: the playbook for monetary policy decisions,”

Macklem, T. (2001) “A New Measure of Core Inflation,” *Bank of Canada Review*

Martínez-García, E., J. Coulter, and V. Grossman (2021) “Fed’s New Inflation Targeting Policy Seeks to Maintain Well-Anchored Expectations,” *Dallas Fed Economics*

Mu, J., and G. Voss (forthcoming) “Inflation Targets and Policy Horizons,”

Norges Bank (2022) “Monetary Policy Report,”

Reserve Bank of Australia (2016) “Statement on the Conduct of Monetary Policy,”

——— (2022) “Statement on Monetary Policy ,”

Reserve Bank of New Zealand (2020) “Monetary Policy Handbook,”

——— (2021) “Monetary Policy Committee Remit,”

——— (2022) “Monetary Policy Statement,”

Rudebusch, G., and L. E. O. Svensson (1999) “Policy Rules for Inflation Targeting,” in J. B. Taylor, ed., *Monetary Policy Rules*, chapter 5, University of Chicago Press

Svensson, L. E. O. (1997) “Inflation forecast targeting: Implementing and monitoring inflation targets,” *European Economic Review* 41(6), 1111–1146

——— (1998) “Monetary policy and inflation targeting,” *NBER Reporter* Winter(1), 5–8

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

——— (1999b) “Inflation targeting: Some extensions,” *The Scandinavian Journal of Economics* 101(3), 337–361

——— (2002) “Inflation targeting: Should it be modeled as an instrument rule or a targeting rule?,” *European Economic Review* 46(4), 771–780

Svensson, L. E. O., and M. Woodford (2005) “Implementing Optimal Policy through Inflation-Forecast Targeting,” in *The Inflation-Targeting Debate*, University of Chicago Press

Sveriges Riksbank (2016) “The Riksbank’s inflation target - target variable and interval,” *Riksbank Studies*

——— (2021) “Account of monetary policy,”

Swiss National Bank (2022a) “Monetary Policy Strategy,”

——— (2022b) “Quarterly Bulletin,”

Taylor, J. B. (1993) "Discretion versus policy rules in practice," *Carnegie-Rochester Conference Series on Public Policy* 39, 195–214

Woodford, M. (2004) "Inflation targeting and optimal monetary policy," *Federal Reserve Bank of St. Louis Review* 86, 15–42

——— (2007) *Forecast Targeting as a Monetary Policy Strategy: Policy Rules in Practice*, 13716, NBER Working Paper

Appendix A

Table 5.1: Recursive Policy Horizons for All-Items CPI Inflation

(a) Absolute Convergence				(b) Relative Convergence			
μ	0.25	0.50	0.75	μ	0.90	0.75	0.60
				$\bar{\mu}^{\dagger}$	0.17	0.42	0.67
Mean	31.17	19.47	10.74	Mean	38.90	26.31	12.33
Median	30.00	20.50	7.00	Median	36.00	26.00	10.00
Std. Dev.	8.89	9.18	8.56	Std. Dev	9.64	7.80	7.89
Min	6.00	4.00	1.00	Min	28.00	6.00	4.00
Max	56.00	34.00	30.00	Max	60.00	55.00	36.00
Obs.	70	70	70	Obs.	70	70	70

† Mean tolerance for convergence ($\hat{\pi} \pm \bar{\mu}_{\epsilon}$) conditional on χ

Table 5.2: Recursive Policy Horizons for CPIX Inflation

(a) Absolute Convergence				(b) Relative Convergence			
μ	0.25	0.50	0.75	μ	0.90	0.75	0.60
				$\bar{\mu}^{\dagger}$	0.07	0.18	0.29
Mean	11.49	4.23	2.77	Mean	38.74	23.20	17.74
Median	8.00	1.50	1.00	Median	33.50	21.00	10.00
Std. Dev.	9.74	4.05	3.26	Std. Dev	12.54	16.00	15.68
Min	1.00	1.00	1.00	Min	12.00	5.00	2.00
Max	34.00	19.00	17.00	Max	61.00	61.00	61.00
Obs.	70	70	70	Obs.	70	70	70

† Mean tolerance for convergence ($\hat{\pi} \pm \bar{\mu}_{\epsilon}$) conditional on χ