BANK OF CANADA
BANQUE DU CANADA

Staff Discussion Paper/Document d'analyse du personnel — 2021-13

Last updated: August 13, 2021

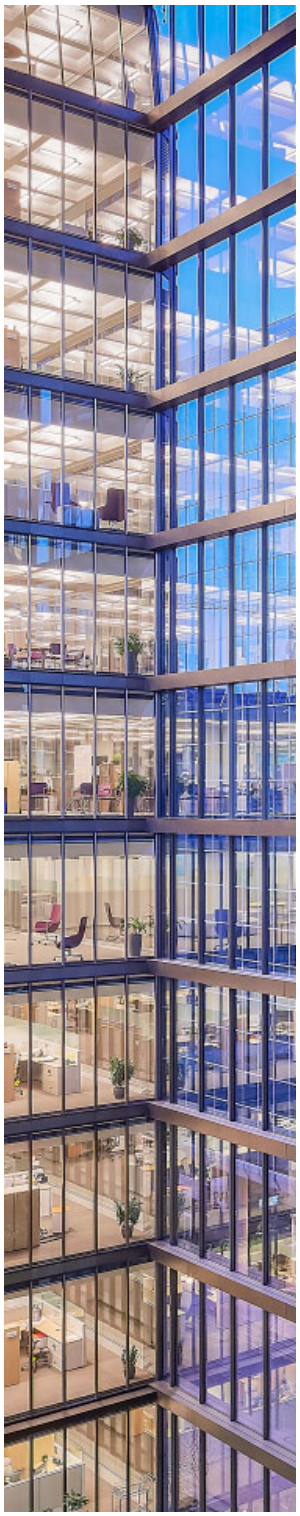

The Bank of Canada's "Horse Race" of Alternative Monetary Policy Frameworks: Some Interim Results from Model Simulations

by José Dorich,[1] Rhys Mendes[1] and Yang Zhang[2]

[1]International Economic Analysis Department
Bank of Canada, Ottawa, Ontario, Canada K1A 0G9

[2]Canadian Economic Analysis Department
Bank of Canada, Ottawa, Ontario, Canada K1A 0G9

jdorich@bankofcanada.ca, rmendes@bankofcanada.ca, yangzhang@bankofcanada.ca

ISSN 1914-0568

Acknowledgements

We are grateful to many Bank of Canada staff who contributed to the work summarized in this paper, including Robert Amano, Thomas Carter, Vivian Chu, Paul Corrigan, Edouard Djeutem, Daniela Hauser, Fanny McKellips, Francis Niyomwungere, Christian Pavilanis, Hassan Raiesi, Abeer Reza, Isabelle Salle, Tudor Schlanger, Jonathan Swarbrick, Joel Wagner, Houji Yao and Jelena Zivanovic. We thank Joe Gagnon, Pierre Fortin and Stephanie Schmitt-Grohe for insightful comments and many participants for their helpful suggestions at the August 2020 workshop, "Towards the 2021 Renewal of the Bank of Canada's Monetary Policy Framework." We also thank Alison Arnot and Colette Stoeber for excellent editorial assistance. The views expressed in this paper are solely those of the authors and may differ from official Bank of Canada views. No responsibility for them should be attributed to the Bank. All simulation results in this paper are obtained based on information up to August 2020.

Abstract

Since 1991, the Bank of Canada has had an inflation-targeting (IT) framework established by a joint agreement between the Bank and the Government of Canada. The framework is reviewed every five years as part of the process for renewing the inflation-control agreement. This discussion paper summarizes some interim results from Bank staff analysis done for the August 2020 workshop, "Towards the 2021 Renewal of the Bank of Canada's Monetary Policy Framework." The Bank will publish updated analysis later in 2021. The core of the current framework—the 2 percent inflation target—has remained unchanged since 1995. This fact reflects its success. Well-anchored inflation expectations contribute to macroeconomic stability while leaving monetary policy with greater flexibility. The 2021 renewal highlights two key challenges facing Canadian monetary policy: (1) the low neutral rate of interest; and (2) the low interest rates associated with a low neutral rate that may encourage excessive risk taking and debt accumulation To address these challenges, Bank staff are running a "horse race" of alternative monetary policy frameworks (i.e., alternatives to the 2 percent IT framework). Their work evaluates these alternatives using a broad range of qualitative and quantitative criteria and focuses on the macroeconomic performance of the alternative frameworks. The interim results we report in this discussion paper suggest overall that no framework dominates on all margins. As a result, the ranking depends on the relative weight placed on different criteria.

Topics: Central bank research; Economic models; Inflation targets; Monetary policy framework; Monetary policy; Monetary policy transmission
JEL codes: E, E2, E27, E3, E4, E5, E52, E58

1. Introduction

Since 1991, the Bank of Canada has had an inflation-targeting (IT) framework established by a joint agreement between the Bank and the Government of Canada.[1] The agreement has been renewed every few years.[2] The most recent agreement was signed in 2016 and runs to the end of 2021. It defines the inflation target as the 2 percent midpoint of a 1 to 3 percent control range for the 12-month rate of change of the total consumer price index (CPI).

The framework is regularly reviewed as part of the process for renewing the inflation-control agreement. These reviews have explored several issues, including:

- the impact of downward nominal wage rigidity and the effective lower bound (ELB) on nominal interest rates
- the appropriate role of financial stability considerations
- the appropriate horizon for returning inflation to target
- the merits of changing the level of the inflation target or adopting price-level targeting

While the framework has evolved, the core of the framework—the 2 percent inflation target—has remained unchanged since 1995. This fact reflects its success. Inflation expectations have become strongly anchored at the target. Such anchoring contributes to macroeconomic stability while leaving monetary policy with greater flexibility to take account of output, employment and developments related to financial stability.

Nevertheless, the IT framework faces important challenges. The main challenge is the low neutral rate of interest, which means that central banks have less room to act in the face of large negative shocks and will need to resort to their expanded set of tools more often. At the same time, the anchoring of inflation expectations has made the relative flatness of the Phillips curve more evident. This, in turn, suggests that inflation will not always provide a clear signal about the maximum sustainable levels of employment and output. With these challenges in mind, the Bank undertook a horse race of alternative monetary policy frameworks (Wilkins 2018).

This discussion paper summarizes interim results from this horse race. These interim results were originally presented at the Bank's August 2020 conference, "Towards the 2021 renewal of the Bank of Canada's Monetary Policy Framework." The results build on work from past reviews

[1] See Amano, Carter and Schembri (2020) and Carter, Mendes and Schembri (2018).
[2] In the 1990s, the frequency of renewals varied. Since the early 2000s, the process has been more regular, with the agreement being renewed every five years.

that looked at changes to the objective of monetary policy. The work for the current review differs from past reviews in its breadth. The Bank is looking at a wide range of alternative frameworks for the 2021 renewal, including:

- average inflation targeting (AIT)
- price-level targeting (PLT)
- an unemployment-inflation dual mandate
- nominal gross domestic product (NGDP) growth and level targeting

Some of these alternatives, such as PLT, imply greater history dependence than conventional IT.[3] In principle, history dependence can lead to better performance of monetary policy in a low neutral rate environment, but it can also lead to greater output volatility if economic agents do not have rational expectations. Other frameworks, such as the employment-inflation dual mandate and NGDP targeting, emphasize the stabilization of a specific real variable more explicitly than IT does. However, IT as practised in Canada and elsewhere is a *flexible* inflation targeting (FIT) framework. That is, the Bank considers the real side of the economy in pursuing its inflation target even though no real objective is identified explicitly in the definition of the framework. So, these alternative frameworks differ from flexible inflation targeting in their explicit identification of a specific real objective. This explicitness would have potential implications for the communication of monetary policy.

The horse race work aims to evaluate these alternative frameworks using a broad range of qualitative and quantitative criteria. These include macroeconomic stability (both price stability and stability of the real economy), financial stability, distributional implications, robustness (to different economic circumstances and different assumptions about private sector behaviour), as well as implications for accountability, communications and credibility. The assessment involves:

- simulations in several different macroeconomic models
- laboratory experiments to assess how well real people understand each framework
- public consultations (see Bank of Canada 2021)

This document summarizes interim results from simulation analysis conducted by Bank staff in three macroeconomic models. The first is the Terms-of-Trade Economic Model (ToTEM), one of the Bank's main policy analysis and projection models. The ToTEM results are then

[3] History dependence means that monetary policy responds to past conditions in addition to current and expected future economic conditions. Section 2 provides further discussion on history dependence.

complemented by analysis in purpose-built models that focus on evaluating the impact of heterogeneity and bounded rationality.

The interim simulation results suggest that no framework dominates on all margins, so the ranking depends ultimately on the relative weight placed on different criteria. Nevertheless, the following notable results emerge:

- The first three frameworks we consider—FIT, AIT and PLT—differ only in the degree of history dependence they embed. We find that their performance depends critically on the importance of the ELB constraint:
    - If the ELB is absent, or if the central bank's extended tool kit can effectively offset its impact, then history dependence does not have significant benefits. Indeed, plausible departures from fully rational expectations cause very history-dependent policies to increase the volatility of the real economy. Thus, FIT tends to dominate AIT and PLT under this situation.
    - If the ELB is an important constraint, then history dependence can be beneficial even in the presence of departures from rational expectations. When the ELB binds, policies that are more history dependent automatically keep the policy rate at the ELB for a longer period. AIT best balances this effect with the additional volatility that history dependence can produce away from the ELB.
- When we broaden the horse race to include frameworks that more explicitly emphasize stabilizing a real variable—the unemployment-inflation dual mandate, NGDP growth targeting and NGDP level targeting—some additional results emerge:
    - Among these additional frameworks, the unemployment-inflation dual mandate outperforms both variants of NGDP targeting in most dimensions.
    - Among all the frameworks, FIT, AIT and the unemployment-inflation dual mandate stand out as the most robust. Notably, however, the dual mandate generates only modest improvements in employment performance over FIT.
    - Frameworks that are more history dependent tend to perform better in large downside scenarios and reduce the frequency of very negative inflation outcomes at the ELB.
    - While each framework has unique strengths and weaknesses, the overall differences in unconditional volatilities of key variables are small relative to the shifts seen over history.
- The analysis in the models with heterogeneous agents and bounded rationality highlights potential costs associated with very high degrees of history dependence:
    - The model with heterogeneous agents relates the cyclical variation in inequality to the cyclical variation in the output gap and the real interest rate. Under

certain circumstances, the very history-dependent frameworks tend to generate more output gap volatility and therefore lead to greater cyclical variation in inequality.[4]

- The model with bounded rationality reinforces the sensitivity of highly history-dependent frameworks to assumptions about expectations formation.

The remainder of this document is organized as follows. Section 2 provides an overview of ToTEM, the benchmark model used in the horse race. Section 3 compares FIT, AIT and PLT. Section 4 expands the horse race to include the unemployment-inflation dual mandate, NGDP growth targeting and NGDP level targeting. Section 5 provides the robustness analysis from other models. Section 6 offers some concluding remarks.

2. Brief overview of ToTEM

ToTEM is a large-scale open economy dynamic stochastic general equilibrium (DSGE) model of the Canadian economy.[5] One of its distinctive properties is that it features significantly more firm- and household-level disaggregation than well-known DSGE models, such as those of Christiano, Eichenbaum and Evans (2005) and Smets and Wouters (2007).

On the firm side, the model features five distinct sectors producing final goods for consumption, residential investment, business investment, government spending and non-commodity exports. The model also includes a separate commodity-producing sector, output from which is mostly exported (as is the case in the data). This elaborate sectoral structure helps the model capture the composition of Canadian gross domestic product (GDP), which is important to accurately evaluate monetary policy frameworks that target the level or growth rate of NGDP or incorporate some role for the output gap.

The firms responsible for producing final goods face nominal rigidities when setting their prices. More specifically, in a given final-good-producing sector, some of the firms re-optimize their prices in a forward-looking but staggered fashion, as in the literature following Calvo (1983), while the other firms set their prices using a rule of thumb (RoT) similar to that in Galí and Gertler (1999).

Estimations of the sector-specific shares for each of these two pricing types find that the share of RoT price setters is relatively high in some sectors. For example, their share in the core consumption sector is estimated, using data from 1995 to 2015, to be about 50 percent. This is

[4] This result is obtained in a model without rule-of-thumb price setters. Ongoing work has expanded analysis to the performance of frameworks in models with some backward-looking behaviour.
[5] See the technical report, Corrigan et al. (2021), for a detailed description of an updated version of ToTEM.

especially important because higher shares are well known to undermine the performance of PLT and other history-dependent regimes that rely more heavily on the expectations channel.[6]

Turning to the household block of ToTEM, we note that the model features three prominent household types differing in terms of the financial markets they have access to and their status as savers or borrowers in those markets.

On the saver side, the model follows Andrés, López-Salido and Nelson (2004) and Chen, Curdia and Ferrero (2012) in assuming that some savers are "restricted" (they can access only long-term debt markets) while others are "unrestricted" (they have access to short- and long-term debt markets). As a result of these two saver types, ToTEM allows short- and long-term interest rates to influence aggregate household spending in distinct ways. This feature of the model could be particularly important in future research evaluating the complementarities between large-scale asset purchases and the frameworks considered here. Taken together, the two saver types—restricted and unrestricted—account for roughly half of all households in the economy.

A single borrower type accounts for most of the remaining households in the economy.[7] Borrowers have been modelled in line with Alpanda and Zubairy (2017). They finance part of their spending using long-term loans secured from saver households. When doing this, borrowers are assumed to face a collateral constraint under which new loans must be backed by some combination of new housing investment and home equity. The first of these components in the collateral constraint is meant to capture mortgages, and the second aims to capture home equity lines of credit. Since these two funding sources account for more than 80 percent of total household debt in Canada, the model offers some insights into a given regime's likely implications for household indebtedness and financial stability.

Turning to the labour market, ToTEM follows most of the DSGE literature in assuming that workers enjoy some degree of wage-setting power but are subject to nominal rigidities similar to those faced by price setters. In particular, a fraction of the economy's wage setters re-optimize their wages on a forward-looking but staggered basis, while the remainder follow an RoT. Shares of the two types have been estimated, with a sizable share of RoT types. This feature of the model is another quantitatively important dimension in which ToTEM departs from the textbook assumption of fully rational, forward-looking behaviour. As mentioned

[6] Sensitivity analysis of various RoT price and wage setting behaviours suggest that this assumption is of great importance for regimes that depend on the expectations channel such as PLT.

[7] The remaining households in the economy represent a "current income" type. Households of this type are assumed not to have access to financial markets and thus simply consume their income on a period-by-period basis. This type accounts for less than 5 percent of all households.

earlier, this is especially important when evaluating the performance of monetary policy frameworks that rely heavily on expectational mechanisms.

Given the structure of the wage-setting process, coupled with the labour demand profile arising on the firm side, the model pins down both the aggregate wage and total number of hours worked in the economy. However, the model does not explicitly include unemployment, which represents a key input in one of our specifications of a dual mandate framework. We therefore estimate a simple equation relating unemployment to hours worked and add it to the model. As expected, this equation suggests a negative relationship between these two variables. Though the equation is reduced-form, its residuals presumably include factors such as movements in labour force participation and the efficiency of the matching process between job seekers and vacancies. Moreover, the unemployment gap in the model is not proportional to the output gap, as would be implied by a simple Okun's law.

Turning finally to the model's policy block, we note that the baseline specification of monetary policy in ToTEM involves a simple rule under which the interest rate is set as a linear function of the previous period's interest rate, the output gap and the deviation of expected inflation over the next four quarters from the central bank's inflation target. On the fiscal side, the government uses a combination of distortionary taxes and bond issuance to finance government spending and transfers. The policy rules governing these expenditures have been estimated and imply government expenditures are countercyclical.

3. ToTEM results: FIT, AIT and PLT

The first three frameworks we consider differ only in the degree of history dependence embedded in the targeted price variable. FIT, AIT and PLT all involve targeting a variable of the form:

$$z_t(N) \equiv \sum_{j=0}^{N} \pi_{t-j}, \tag{1}$$

where π_t is the month-over-month inflation rate. When $N = 12$, z_t corresponds to the year-over-year inflation rate that is targeted in conventional FIT frameworks. Since the price level is the cumulation of all past price changes, in the limit as $N \to \infty$, z_t converges to the price level. Thus, PLT can be viewed as a limiting case. Intermediate cases ($12 < N < \infty$) correspond to AIT—targeting a multi-year average of inflation.

Larger values of N entail reacting to longer histories of inflation rates. A central bank that has a target variable that is very history dependent will attempt to "make up" for misses further in the past. For this reason, history-dependent policies are sometimes referred to as "make up" strategies.

FIT, AIT and PLT can be nested in a policy rule of the form:

$$i_t = 0.85 i_{t-1} + (1 - 0.85)\left\{i^* + \gamma\left(\frac{12}{N}\right)\sum_{j=0}^{N}\left(\pi_{t-j} - \bar{\pi}\right) + \alpha\tilde{x}_t\right\}, \quad (2)$$

where i_t is the nominal policy interest rate, $\tilde{x}_t$ is the output gap, i^* is the neutral level of the policy interest rate, and $\bar{\pi}$ is the inflation target expressed at a monthly rate under FIT and AIT and the trend inflation rate under PLT.[8,9]

For concreteness, we consider three cases:

$$FIT: i_t = 0.85 i_{t-1} + (1 - 0.85)\left\{i^* + \gamma_{yy}\left(\pi_t^{yy} - \bar{\pi}^a\right) + \alpha_{yy}\tilde{x}_t\right\} \quad (3)$$

$$AIT: i_t = 0.85 i_{t-1} + (1 - 0.85)\left\{i^* + \gamma_{3y}\left(\pi_t^{3y} - \bar{\pi}^a\right) + \alpha_{3y}\tilde{x}_t\right\} \quad (4)$$

$$PLT: i_t = 0.85 i_{t-1} + (1 - 0.85)\left\{i^* + \gamma_p(p_t - \bar{p}_t) + \alpha_p\tilde{x}_t\right\}, \quad (5)$$

where $\pi_t^{yy} \equiv z_t(12)$ is the year-over-year inflation rate, $\pi_t^{3y} \equiv (1/3)z_t(36)$ is the three-year average rate of inflation expressed at an annual rate and $\bar{\pi}^a$ is a constant target set at 2 percent in our simulations. The price-level target has a deterministic trend: $\bar{p}_t = \bar{p}_{t-1} + \bar{\pi}^a$. The parameters γ and α are chosen to minimize a loss function, as discussed below. Note also that we considered AIT variants between two and five years. We report results only for the three-year variant as it outperformed the others.

The benefits of history dependence

To understand the benefits of history dependence in forward-looking environments, it is useful to look at the example of PLT. The traditional literature on the choice between FIT and PLT framed the decision as a trade-off between long-run price-level uncertainty and short-run variability of inflation and output. This conclusion arose naturally in models with a limited role for forward-looking behaviour: after a positive shock to the price-level, a central bank would have to induce a contraction in demand to force inflation below trend to return the price level

[8] We have in mind PLT regimes in which there can be a deterministic trend in the targeted path for the price level. In this report, we focus on contemporaneous inflation only. Ongoing work has expanded analysis to the performance of forecast rules.

[9] The presence of the output gap is consistent with an optimal simple rule that is obtained by minimizing a model-consistent concept of welfare loss. Having a measure of real economic slack is beneficial for evaluating monetary policy rules that could respond appropriately to shocks that potentially cause economic activity and inflation to move in opposite directions.

to its target. Given these results, the choice between FIT and PLT is reduced to an assessment of the relative importance of long-run price-level certainty.

The traditional literature, however, largely ignored the role of endogenous expectations. Svensson (1999) demonstrated that this omission was not innocuous. Under plausible conditions, introducing forward-looking behaviour could allow for a "free lunch": PLT could generate both lower long-run price-level uncertainty and lower short-run variability in inflation and output.

Svensson's "free lunch" result is a direct consequence of the effect of PLT on the behaviour of inflation expectations. Under PLT, the expectation that policy-makers will undo the effects of shocks to the price level discourages firms from changing prices as dramatically as they would under a regime that accommodated shocks. This adjustment in inflation expectations is particularly important when the economy hits the ELB, since higher expected inflation helps reduce real interest rates. The fact that bygones are not bygones in a PLT regime renders policy history dependent in a manner that mimics the type of history dependence that characterizes optimal policy in forward-looking models (Woodford 2003).

Of course, for this history dependence to induce the beneficial automatic adjustment of expectations, economic agents must understand the nature of the regime and they must believe that the central bank's commitment to the regime is credible. If these conditions are not satisfied, PLT can lead to greater short-run volatility in inflation and output. AIT offers an intermediate degree of history dependence because central banks only need to make up for misses over a finite averaging window. Shocks to the price level eventually drop out of the averaging window, so, under AIT, bygones *eventually* are bygones.

Key assumptions

For the analysis in this section and the next one, we use ToTEM. As explained above, ToTEM has multiple sectors, so it has multiple price levels and rates of inflation. We assume that FIT, AIT and PLT are all based on consumer prices in the model. The Bank's actual FIT framework involves targeting the *total* CPI rate of inflation. However, the forward-looking nature of real-world monetary policy means that the Bank can "look through" transitory volatility in total CPI inflation. We use core inflation in the rules in equations (3) to (5) to capture the idea that the Bank can look through temporary volatility.[10] Note also that the model contains a unique

[10] An alternative would be to use forward-looking rules based on the total CPI. Such rules would look through temporary volatility in a model-consistent manner. This is being considered in ongoing work.

concept of core, so we do not have to choose among the Bank's three empirical measures of core inflation.[11]

We consider the performance of the rules in equations (3) to (5) for different values of the γ and α parameters, but we fix the smoothing parameter at 0.85. This value is broadly in line with estimates of simple monetary policy rules for Canada and other jurisdictions. We fix this parameter at a constant value in order to distinctly compare different frameworks. Differences in the smoothing parameter could otherwise confound differences in the degree of history dependence embedded in the target variables.

Assumptions about the effectiveness of the extended monetary policy (EMP) tool kit (forward guidance, quantitative easing, etc.) at the ELB are important for our evaluation of alternative frameworks. In work for the 2016 renewal, Bank staff explicitly modelled the effects of the EMP tools. That approach is being implemented in ongoing work but is not reported in this paper. Rather, we deal with the EMP tool kit by considering two polar cases:

(1) The policy interest rate is subject to an occasionally binding ELB at 0.25 percent, and no other monetary policy tools are available at the ELB. One can think of this case as approximating a situation in which the EMP tool kit is completely ineffective.
(2) The policy interest rate is unconstrained. Negative values of the policy rate in the simulation can be interpreted as representing the shadow policy interest rate (a measure of the stance of monetary policy taking account of EMP tools).[12] This can be viewed as approximating a situation in which the EMP tool kit can perfectly substitute for desired reductions in the policy interest rate below the ELB.

These two cases represent extreme assumptions about the effectiveness of the EMP tool kit. Reality lies somewhere in between, with the EMP tool kit only able to partially substitute for desired reductions in the policy rate below the ELB. By focusing on the extreme cases, our approach yields information on the robustness of the alternative monetary policy frameworks to the availability and effectiveness of EMP tools.

ToTEM results for FIT, AIT and PLT

We begin by assuming that the central bank aims to minimize the volatility of both inflation and the output gap. As a first step, we can be agnostic about the precise loss function and instead compare the efficient policy frontiers for each of the frameworks. We perform the analysis using stochastic simulations for each regime, allowing for different coefficient values in

[11] ToTEM is estimated using core inflation defined as an average of the three preferred measures of core CPI.
[12] We have in mind the concept of a shadow rate in the sense of Black (1995) as applied by Wu and Xia (2016).

respective policy rules. The shocks are drawn from the estimated historical distribution of shocks in the model.

Chart 1 plots efficient frontiers for each framework both with and without the ELB. The vertical axis shows the mean squared errors (MSEs) of the output gap, and the horizontal axis shows the MSEs of annualized inflation. The frontiers are derived by searching over the γ and α parameters and plotting the lowest achievable pairs of the unconditional MSEs of the output and inflation gaps.[13] The frontiers represent the efficient set in the sense that the variability of inflation can only be reduced further if the variability in the output gap is increased, or vice versa.

- In the absence of the ELB (solid lines in **Chart 1**), standard FIT performs very well. Given our interpretation of the "no ELB" case, this suggests that FIT would do well in a situation in which the EMP tool kit is highly effective. In this type of situation, the additional history dependence inherent in AIT and PLT frameworks does not improve inflation and output gap volatility.
- With an occasionally binding ELB (dashed lines in **Chart 1**), FIT is always dominated by AIT and it is dominated by PLT for parameterizations that lead to low inflation volatility. The weaker performance of FIT in this case reflects the fact that history dependence is more important when the ELB binds. In this situation, policies that are more history dependent automatically keep the policy rate at the ELB for a longer period. This provides additional stimulus that mitigates the effect of the ELB.

[13] The frontiers imply in some cases very volatile interest rates.

Chart 1: Policy frontiers

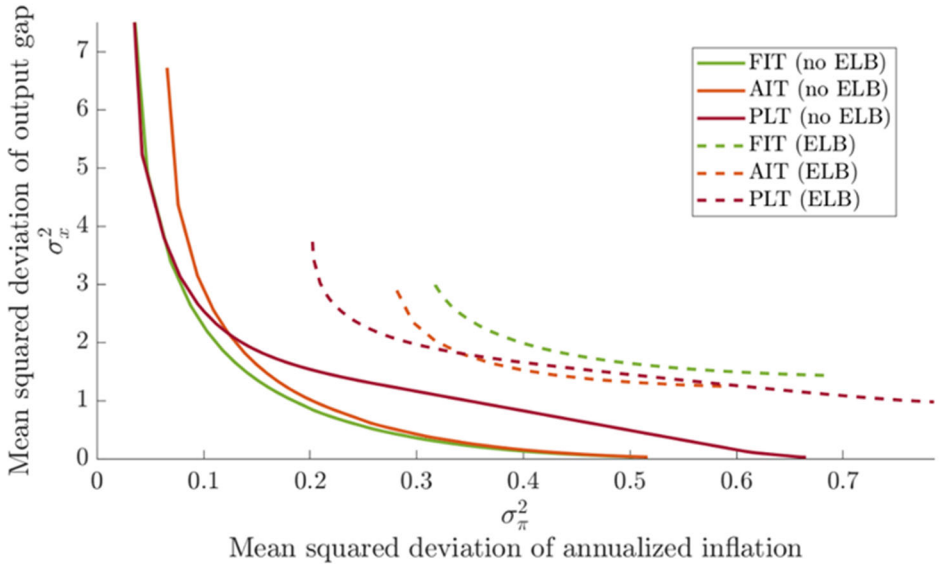

Note: FIT is flexible inflation targeting; AIT is average inflation targeting; PLT is price-level targeting; and ELB is the effective lower bound.

These results from ToTEM contrast with the dominance of PLT in the canonical New Keynesian model. The difference stems from the high proportion of RoT wage and price setters in ToTEM. This is qualitatively consistent with the findings of Amano et al. (2020). They show that in the presence of RoT firms, highly history-dependent frameworks generate a worse inflation-output trade-off. Thus, RoT price setters reduce the optimal degree of history dependence.[14] Amano et al. (2020) also show that, all else equal, a higher degree of history dependence is optimal when there is an occasionally binding ELB constraint, consistent with the results in **Chart 1**.

We next evaluate the frameworks using a loss function of the form:

$$L = E[(\pi_t^a - \bar{\pi}^a)^2 + \lambda_x \tilde{x}_t^2 + 0.5(i_t - i_{t-1})^2], \tag{6}$$

where $E[\cdot]$ is the unconditional expectation operator, π_t^a is the annualized quarter-over-quarter inflation rate, $\tilde{x}_t$ and i_t are the output gap and nominal annualized quarter-over-quarter interest rate as before, $\bar{\pi}^a$ is the inflation target and λ_x is the relative weight on the

[14] Our results assume that the proportion of RoT price setters is a structural feature of the economy. This implies that it is invariant to changes in the monetary policy framework. In reality, firms make a choice about how to behave. Changes in the monetary policy framework could lead to endogenous changes in the price-setting behaviour of firms. Bank staff are looking at this question in a model that allows firms to choose whether or not to follow a simple RoT.

output gap. Much of the Bank's past work on the monetary policy framework has assumed $\lambda_x = 1$. One possible in-model interpretation of a "dual mandate" is a larger value of λ_x. To look at what type of framework would perform well for different values of λ_x we select the γ and α parameters in equations (3) to (5) to minimize the loss function in (6). **Charts 2** and **3** plot the value of the minimized loss function for each framework for values of λ_x between 0 and 2.

Chart 2 focuses on the case in which the ELB is not a constraint (or the EMP tool kit is very effective). In this case, FIT generates the lowest loss for $\lambda_x > 0.26$. This suggests that if the EMP tool kit is effective, an inflation target would perform best in an environment where the central bank significantly emphasizes stabilizing the real economy. The policies that are more history dependent (AIT and PLT) perform better if the central bank places little weight on stabilizing the output gap. This, again, is due to the presence of RoT wage and price setters. The RoT behaviour worsens the inflation-output trade-off generated by the more history-dependent policies. The inferior trade-off has a greater adverse effect on losses when the central bank cares about variability in both inflation and the output gap. That said, the differences between FIT and AIT are quantitatively small regardless of the value of λ_x.

Chart 2: Loss values with no effective lower bound (ELB)

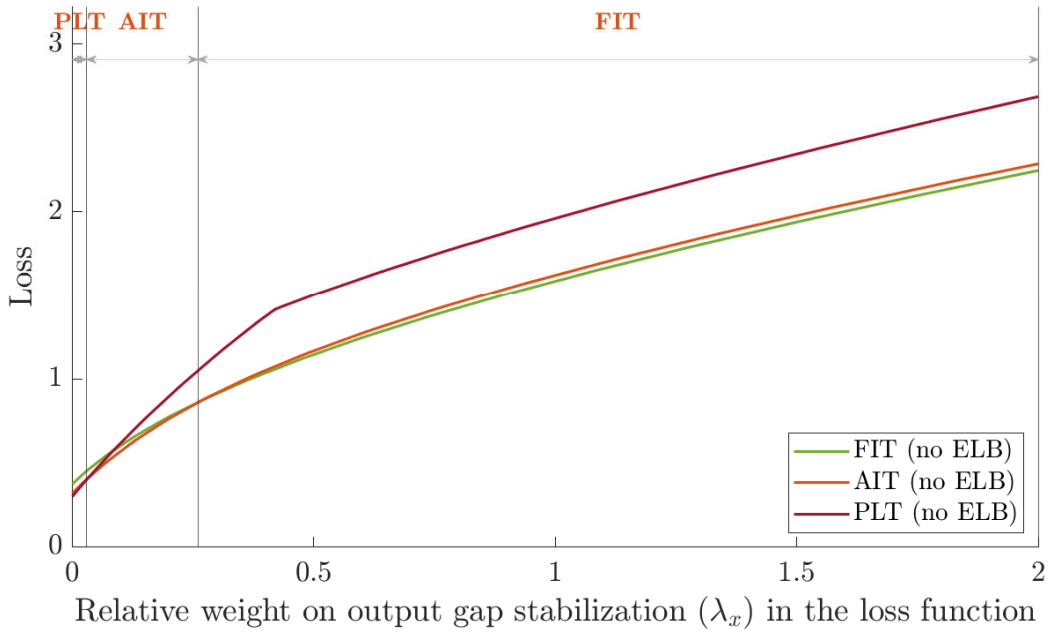

Note: FIT is flexible inflation targeting; AIT is average inflation targeting; PLT is price-level targeting.

Chart 3 shows the losses in the case with an occasionally binding ELB. With the ELB, history dependence yields greater benefits. FIT is now dominated by AIT and PLT for all values of λ_x. AIT yields the lowest loss for a wide range of λ_x values. In other words, if the EMP tool kit is ineffective, an intermediate degree of history dependence helps to stabilize both inflation and the real economy. This, again, reflects the fact that regimes that are more history dependent automatically provide more stimulus at the ELB.

Overall, the results in this section suggest that PLT is dominated by alternatives that are less history dependent regardless of how much relative weight the central bank places on stabilizing the real economy.[15] The horse race between FIT and AIT depends on the assumed effectiveness of the EMP tool kit. If EMP tools are effective at compensating for the ELB, our results suggest that FIT is marginally superior to AIT. In contrast, AIT has a somewhat larger lead over FIT if EMP tools are assumed to be ineffective. Thus, AIT might be regarded as more robust to the effectiveness of the EMP tool kit given that its relative performance is good in both cases.

Chart 3: Loss values with occasionally binding effective lower bound (ELB)

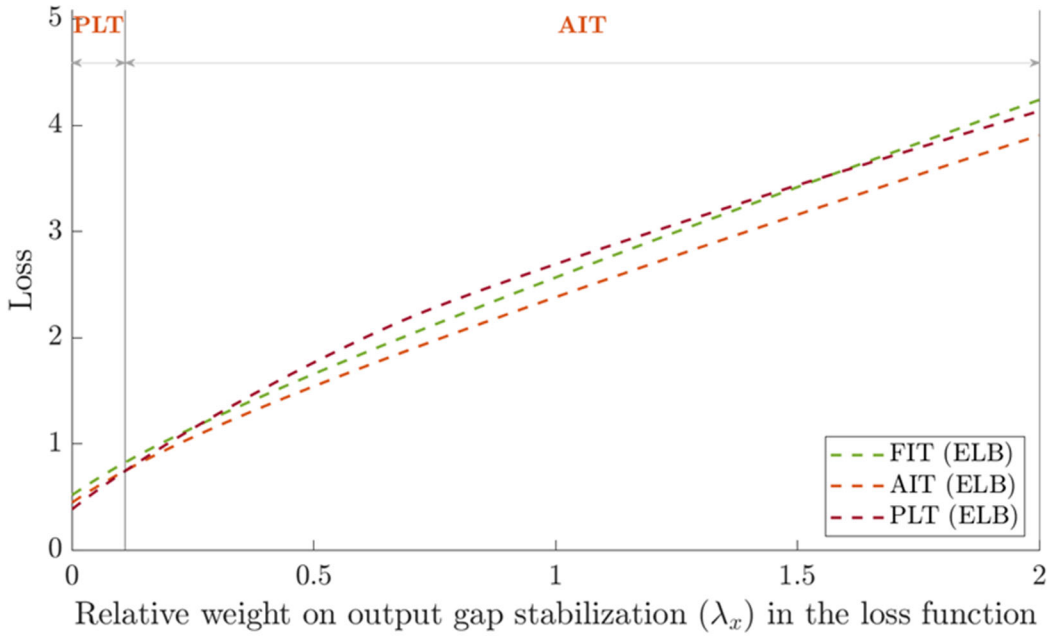

Note: FIT is flexible inflation targeting; AIT is average inflation targeting; PLT is price-level targeting.

[15] This conclusion may be affected by our focus on rules in which the target variable enters contemporaneously. Past research suggests that the performance of PLT can improve when the policy rule is allowed to be forward-looking. See, for example, Coletti, Lalonde and Muir (2008) and Smets (2003).

4. ToTEM results: Expanding the horse race

The frameworks considered in the previous section all involved targeting some derivative of consumer prices. This made it reasonable to evaluate them using a simple ad hoc loss function based on consumer price inflation and the output gap. In this section, we consider a wider set of frameworks, bringing NGDP growth targeting, NGDP level targeting and an unemployment-inflation dual mandate into the horse race. The ad hoc loss function used in the previous section penalizes these frameworks because they target variables that do not enter the loss function. This suggests we need a different approach to run a fair horse race with this broader set of frameworks.

The expanded horse race

A natural approach would be to use a model-consistent social welfare function based on households' utility to evaluate the alternative frameworks. However, two issues arise in doing this:

(1) Significant technical obstacles exist in evaluating welfare in ToTEM with an occasionally binding ELB.
(2) Some features of the model, such as Calvo (1983) pricing, are useful for characterizing macroeconomic dynamics in a tractable way but are based on micro foundations that have unrealistic implications for welfare.

Nevertheless, we continue to work on computing welfare in ToTEM. For the time being, however, we take an approach that remains agnostic about the social welfare function. We proceed as follows:

- We characterize each framework using a regime-specific loss function delegated to the central bank and a simple policy rule (**Table 1**). We use the delegated loss function to choose the parameters of the rule, but not to evaluate the framework.
- We evaluate the frameworks using volatilities of several key economic variables. We do not explicitly assign weights to these volatilities in order to compute a loss function, but instead look for alternative frameworks that stabilize a broad range of variables.

Each of the three new frameworks introduced in this section are described below.

Table 1: Regime-specific delegated loss functions and simple rules

Framework	Loss specification	Interest rate rules*
Flexible inflation targeting	$L^{FIT} = \left(\pi_t^{yy} - \bar{\pi}^a\right)^2 + (\widetilde{x_t})^2 + 0.5\,(\Delta i_t)^2$	$i_t = i^* + \gamma_{yy}\left(\pi_t^{yy} - \bar{\pi}^a\right) + \alpha_{yy}\widetilde{x_t}$
Average inflation targeting	$L^{AIT} = \left(\pi_t^{3y} - \bar{\pi}^a\right)^2 + (\widetilde{x_t})^2 + 0.5\,(\Delta i_t)^2$	$i_t = i^* + \gamma_{3y}\left(\pi_t^{3y} - \bar{\pi}^a\right) + \alpha_{3y}\widetilde{x_t}$
Price-level targeting	$L^{PLT} = (p_t - \bar{p}_t)^2 + (\widetilde{x_t})^2 + 0.5\,(\Delta i_t)^2$	$i_t = i^* + \gamma_p(p_t - \bar{p}_t) + \alpha_p\widetilde{x_t}$
Nominal-GDP level targeting	$L^{NGDPL} = \begin{Bmatrix} (y_t + p_{GDP,t}) \\ -(\bar{y}_t + \overline{p_{GDP,t}}) \end{Bmatrix}^2 + 0.5\,(\Delta i_t)^2$	$i_t = i^* + \delta_{ngdpl}\begin{bmatrix} (y_t + p_{GDP,t}) \\ -(\bar{y}_t + \overline{p_{GDP,t}}) \end{bmatrix}$
Nominal-GDP growth targeting	$L^{NGDPG} = \begin{Bmatrix} (\Delta y_t^{yy} + \Delta p_{GDP,t}^{yy}) \\ -(\overline{\Delta y_t}^{yy} + \overline{\Delta p_{GDP,t}}^{yy}) \end{Bmatrix}^2 + 0.5(\Delta i_t)^2$	$i_t = i^* + \delta_{ngdpg}\begin{bmatrix} (\Delta y_t^{yy} + \Delta p_{GDP,t}^{yy}) \\ -(\overline{\Delta y_t}^{yy} + \overline{\Delta p_{GDP,t}}^{yy}) \end{bmatrix}$
Unemployment-inflation dual mandate	$L^{DM} = \left(\pi_t^{yy} - \bar{\pi}^a\right)^2 + (\widetilde{u_t})^2 + 0.5\,(\Delta i_t)^2$	$i_t = i^* + \gamma_{yy}\left(\pi_t^{yy} - \bar{\pi}^\alpha\right) + \alpha_u\widetilde{u_t}$

*In the simulations, all the frameworks considered assume a common smoothing parameter of 0.85, similar to equations (3) to (5). For simulations with an effective lower bound, the components of the loss functions are expressed in terms of mean squared deviations instead of variances.

Unemployment-inflation dual mandate*:* In the previous section, we discussed the possibility that a dual mandate could be interpreted as a larger weight on a real variable in the loss function. Real-world examples of dual mandates, such as the frameworks of the Federal Reserve System of the United States and the Reserve Bank of New Zealand, usually include an employment or unemployment objective. This type of dual mandate creates a potentially larger role for labour market conditions in the conduct of monetary policy. We model such a regime in much the same way as we did with FIT, except we replace the output gap measure in the loss function and interest rate rule with an unemployment gap, $\widetilde{u_t}$, measuring the difference between the actual and natural rates of unemployment.

Nominal-GDP level targeting*:* NGDP level targeting is modelled as stabilizing the sum of the logarithm of real GDP, y_t, and the logarithm of the GDP deflator, $p_{GDP,t}$, around a deterministic trend. This framework has received renewed attention because of its potential to address some of the challenges of the current environment (see, e.g., Ambler 2020). For instance, like PLT, NGDP level targeting features a high degree of history dependence, which may improve its performance at the ELB. It may also have some appealing implications for financial stability.

Consider a productivity shock that raises real GDP and reduces prices. Under NGDP level targeting, the impetus to ease monetary policy stemming from lower prices is counterbalanced by the increase in output. In contrast, under FIT, inflation and the output gap would both call for lower rates after a positive productivity shock, potentially increasing the incentives for risk taking and debt accumulation.

However, NGDP level targeting has important disadvantages too. For example, if changes in trend real GDP growth are not offset by changes in the target path for NGDP, then trend inflation will be forced to adjust. This would reduce agents' level of certainty about future inflation when making long-term decisions and would likely lead to less well-anchored inflation expectations. In addition, the weights on the price and real variables in the reaction function are constrained to be equal under NGDP level targeting.[16]

Nominal-GDP growth targeting*:* We model NGDP growth targeting as stabilizing the sum of the year-over-year real GDP growth, Δy_t^{yy} and the year-over-year rate of change of the GDP deflator, $\Delta p_{GDP,t}^{yy}$. This framework differs from FIT in two important dimensions. First, it incorporates the rate of change of the GDP deflator rather than CPI inflation. Second, it assigns an explicit role to real economic growth in the determination of interest rates. Unlike all the other frameworks, the relevant real variable enters as a growth rate rather than a level. This distinction will prove to have an important impact on the simulated performance of this framework.

Unconditional analysis in ToTEM

We begin by evaluating the relative performance of the frameworks in terms of how well they stabilize key aggregate variables. When doing so, we entertain two polar cases that differ in their treatments of the ELB on nominal interest rates and the nature of the extended tool kit assumed available to the central bank. As in the previous section, our first case accounts for the ELB but assumes that the extended tool kit cannot be deployed or is ineffective. In contrast, our second case abstracts from the ELB. We interpret this case as a situation in which the extended tool kit can perfectly substitute for conventional monetary stimulus, thus allowing the central bank to achieve the same outcomes as without the ELB. For this reason, we interpret negative rates in this case as representing a shadow rate. Together, these two cases represent natural benchmarks. They provide upper and lower bounds on a given framework's performance in a

[16] An additional limitation on NGDP level targeting is due to delayed information (as the nominal GDP measure is only available quarterly), and there is the potential for large revisions over time.

more realistic scenario in which the central bank has access to an extended tool kit of some form, but the instruments therein cannot fully offset the impact of the ELB.[17]

Our results for the first case are reported in **Table 2**, which presents the unconditional standard deviations of the following key variables under each framework: total CPI inflation (year-over-year), the output gap, real GDP growth (year-over-year), the unemployment rate, the first difference of the nominal interest rate and real household debt growth.[18] Several interesting results stand out.

The first is that no one framework strictly dominates the others in the sense that it is better able to stabilize all the variables in question. However, several frameworks offer notably narrow advantages. For example, NGDP growth targeting is strictly dominated with respect to all variables except for real household debt growth, and its relative advantage in this one aspect is quantitatively small. As a result, it would be favoured only under a ranking system that places nearly singular weight on the stabilization of household debt growth.

Table 2: Unconditional standard deviations (%), with effective lower bound

Policy rule	Total CPI inflation (Y/Y)*	Output gap	First difference of interest rate	Unemployment rate	Real GDP growth (Y/Y)	Real household debt growth (Y/Y)
Flexible inflation targeting	0.67	1.25	0.69	0.89	1.94	5.69
Average inflation targeting	0.68	1.20	0.71	0.87	1.94	5.69
Unemployment-inflation dual mandate	0.70	1.42	0.57	0.78	2.07	5.68
Price-level targeting	0.56	1.56	0.78	0.99	2.19	5.83
Nominal-GDP level targeting	0.84	1.58	1.21	0.99	2.01	5.68
Nominal-GDP growth targeting	1.18	2.27	1.22	1.34	2.10	5.50

*Y/Y refers to year-over-year.

Similarly, PLT emerges as the framework that best stabilizes CPI inflation. Moreover, as shown in **Chart 4**, it is the regime that minimizes the likelihood of very low inflation outcomes.

[17] A more explicit analysis accounting for the likely availability and effectiveness of these instruments in the context of the horse race will be provided in future work.

[18] Future work will report on an expanded set of moments, including the mean outcome of key macroeconomic variables.

However, PLT is strictly dominated by FIT, AIT and the dual mandate with respect to all other variables in the table. This suggests that PLT would only be favoured under a ranking system that places very high weight on inflation stabilization. At the same time, FIT, AIT and the dual mandate all outperform NGDP level targeting with respect to all variables except for real household debt growth, where the margin in question is quantitatively small. As a result, it is fair to assume that most ranking systems would place NGDP level targeting somewhere behind FIT, AIT and the dual mandate.

FIT, AIT and the dual mandate thus generally emerge as the most robust of the frameworks. A key corollary is that an intermediate degree of history dependence best strikes a reasonable balance between:

- the benefits that history dependence has to offer in terms of its stabilizing effects on expectations at the ELB
- the costs stemming from its destabilizing interactions with RoT agents

That said, we note the quantitative differences in unconditional volatilities across the frameworks under consideration are small relative to historical shifts. In particular, substantial declines in macroeconomic volatility occurred after we adopted FIT in Canada. For example, the standard deviation of inflation fell from 3.1 percent in the 1970s and 1980s to 0.8 percent in the 1995 to the first quarter of 2020.[19] Similarly, the standard deviation of real output growth declined from 3.9 percent to 2.6 percent over that period. In contrast, most of the differences in **Table 2** are much smaller.

[19] We omit the first few years of FIT because the inflation target was declining over that period.

Chart 4: Distribution of total inflation under different policy regimes

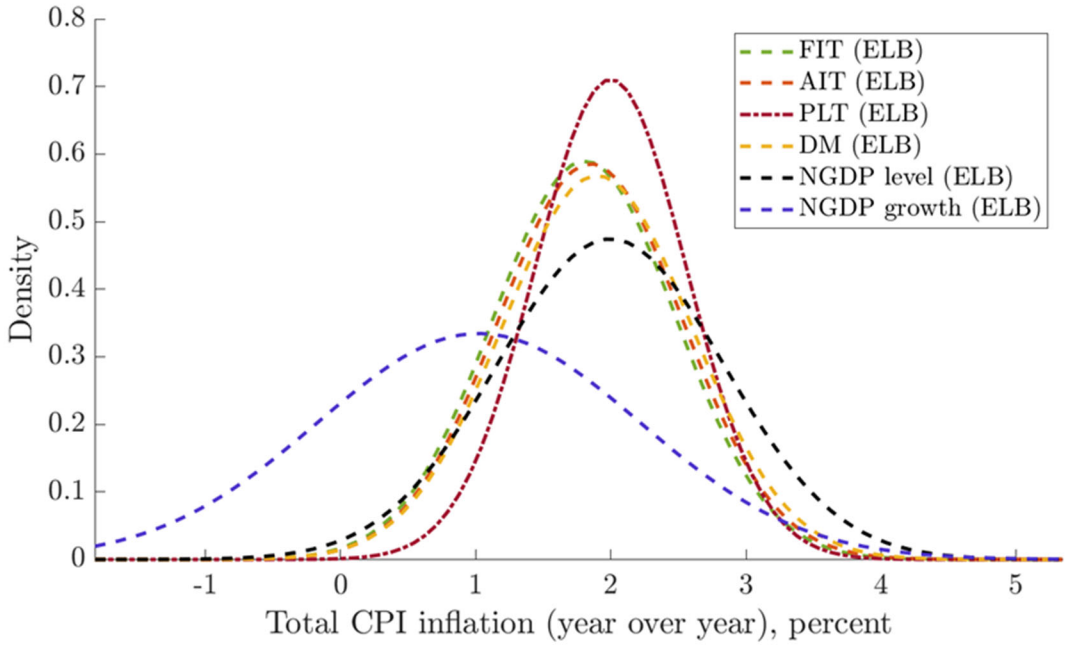

Note: FIT is flexible inflation targeting; AIT is average inflation targeting; PLT is price-level targeting; DM is unemployment-inflation dual mandate; NGDP level is nominal-GDP level targeting; NGDP growth is nominal-GDP growth targeting; and ELB is the effective lower bound.

Table 3 presents our results for the case in which we abstract from the ELB. As in the previous case, no one framework dominates the others with respect to all variables we consider. Moreover, many of the general patterns from our previous case continue to hold in some form. For example, NGDP growth targeting now emerges as the framework that best stabilizes *both* real GDP growth and real household debt growth. However, its advantages relative to FIT and AIT are small and achieved at the cost of considerable volatility in the output gap and large period-to-period changes in the interest rate. As a result, most ranking systems would likely place NGDP growth targeting some distance behind FIT and AIT.

As for PLT, it continues to offer the highest degree of inflation stabilization and now has the additional advantage that it is second-best in terms of its ability to economize on period-to-period changes in the interest rate. However, together with NGDP level targeting, PLT is one of the worst performers in terms of stabilizing all real variables presented in the table. As explained earlier, this relatively poor performance of the two regimes with the highest degree of history dependence mainly reflects the prevalence of RoT behaviour in the economy and the limits that this behaviour places on the strength of expectational transmission mechanisms.

Table 3: Unconditional standard deviations (%), without effective lower bound						
Policy rule	**Total CPI inflation (Y/Y)***	**Output gap**	**First difference of interest rate**	**Unemployment rate**	**Real GDP growth (Y/Y)**	**Real household debt growth (Y/Y)**
Flexible inflation targeting	0.68	0.88	0.87	0.78	1.73	5.64
Average inflation targeting	0.60	0.79	0.79	0.78	1.64	5.55
Unemployment-inflation dual mandate	0.65	1.30	0.63	0.67	2.01	5.66
Price-level targeting	0.51	1.51	0.83	0.97	2.10	5.83
Nominal-GDP level targeting	0.81	1.55	1.61	0.97	2.01	5.77
Nominal-GDP growth targeting	0.75	1.40	2.17	0.91	1.67	5.59

*Y/Y refers to year-over-year.

For these reasons, FIT, AIT and the dual mandate also emerge from our second case as the most robust of the frameworks under consideration. Moreover, their relative merits tend to be the same as in the previous case. For example, under both cases, all three of these regimes perform similarly in terms of stabilizing CPI inflation and real household debt growth. This makes the choice among these regimes dependent principally on the weights that one attaches to stabilizing other variables, such as the output and unemployment gaps, real GDP growth and period-to-period changes in the interest rate. In particular, FIT and AIT tend to do better than the dual mandate in stabilizing the output gap and real GDP growth, while the dual mandate does better in stabilizing unemployment and the variations in the interest rate.

These qualitative patterns are common to both the ELB and no ELB cases. These cases have been designed to bracket a given regime's performance in a situation that takes the central bank's extended tool kit into account more explicitly. As a result, the relative rankings that emerge from our analysis are unlikely to change in the context of an exercise that more formally incorporates the likely availability and effectiveness of the extended tool kit.

Scenario analysis in ToTEM

Though unconditional analyses like the ones presented above are informative, gauging the regimes' conditional performance in large but plausible downside scenarios is also useful. For

this reason, we now turn to an analysis of the regimes' resilience in a scenario in which the economy is hit by a sequence of shocks resembling those that occurred over the 2008–10 period.[20]

Chart 5 shows the responses of key macroeconomic variables to these shocks, expressed in deviations from their steady-state values. Note that the scenario assumes a steady-state nominal neutral rate of 2.75 percent and an ELB of 25 basis points; it also abstracts from the possibility that instruments from the central bank's extended tool kit could be deployed once the nominal rate reaches the ELB.[21]

Performance is roughly similar across all regimes except for NGDP growth targeting, which clearly performs worst. For example, the response of the output gap under NGDP growth targeting reaches about -7 percent at its trough, compared with a range of -4.8 to -4.3 percent under the other regimes. The relatively poor performance of NGDP growth targeting is explained by the considerably shorter period of time for which the ELB binds under this regime: the rule targeting NGDP growth prescribes only 4 quarters at the ELB, while the other regimes imply durations in the range of 9 quarters to 13 quarters. This difference occurs because NGDP growth begins to recover about 4 quarters after the beginning of the simulation, while the level of GDP remains depressed.

Among the other regimes, the differences in performance are very modest. By a small margin, NGDP level targeting seems to best insulate the economy from the shocks in question. Under this regime, trough responses of CPI inflation, the output gap and real GDP growth are the least negative, while the peak response of the unemployment gap is the least positive. This is explained by the relatively more aggressive interest rate response of NGDP level targeting, which allows it to reach the ELB more quickly than the other regimes and achieve larger declines in the long-term rates during the first couple of quarters in the simulation.

[20] We use these shocks to match the historical Canadian data (deviation from respective trends) over a period from the first quarter of 1995 to the fourth quarter of 2015, a sample that we use to estimate the latest version of ToTEM.

[21] Updated simulation results based on the 2021 neutral rate estimate will be published in the inflation target renewal background document and upcoming staff discussion paper on the horse race (Swarbrick and Zhang, forthcoming).

Chart 5: A scenario analysis of policy regime performance in ToTEM

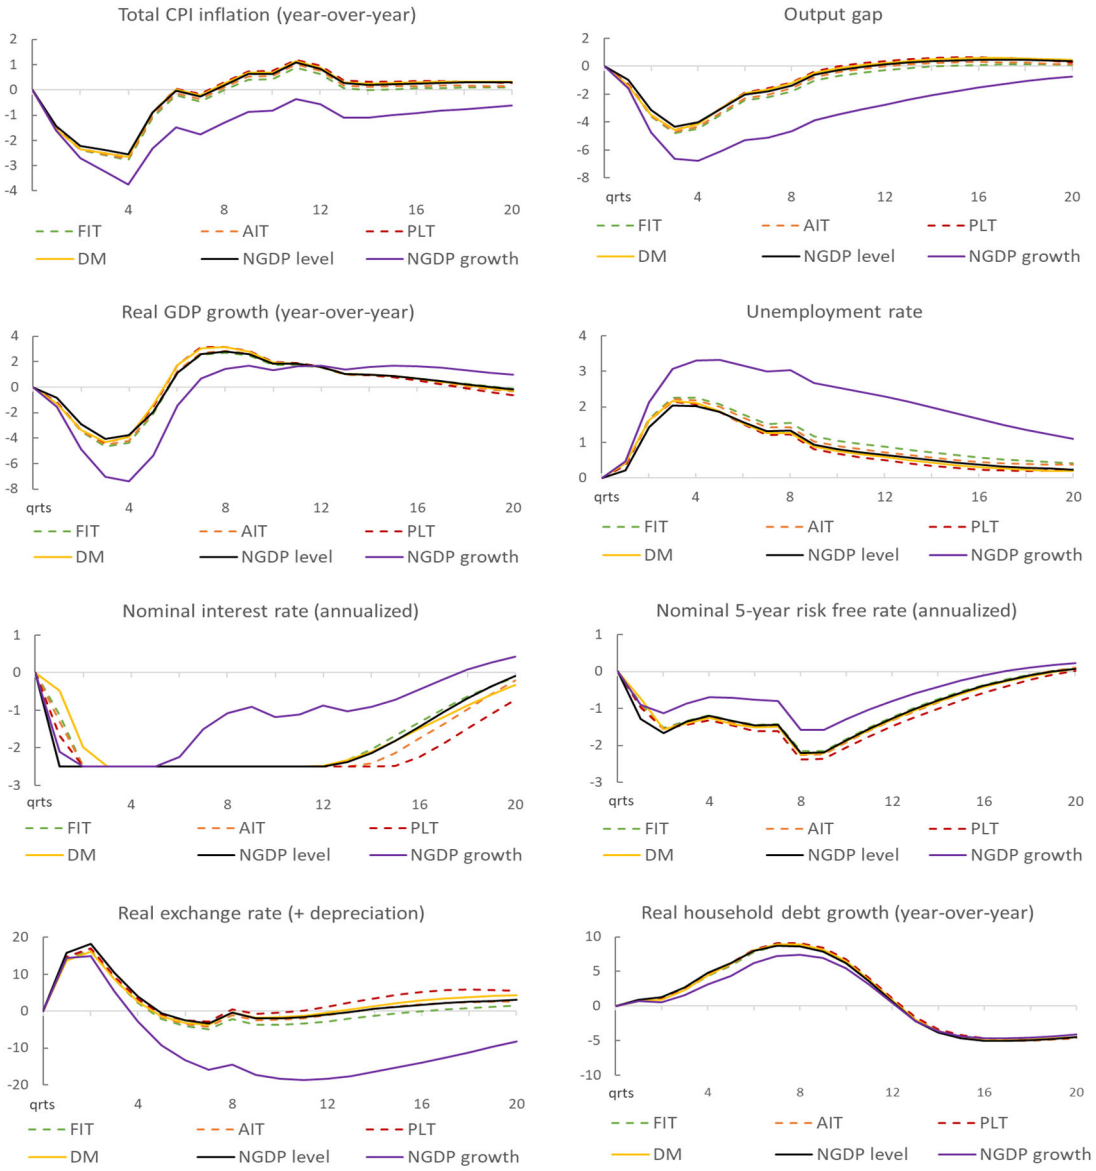

Note: ToTEM is the Terms-of-Trade Economic Model; FIT is flexible inflation targeting; AIT is average inflation targeting; PLT is price-level targeting; DM is unemployment-inflation dual mandate; NGDP level is nominal-GDP level targeting; and NGDP growth is nominal-GDP growth targeting.

Turning finally to the regimes' potential implications for financial stability, we note that the peak responses of real household debt growth fall within a relatively narrow range of 8.7 to 9.1 percentage points, excluding NGDP growth targeting. The relatively limited extent to which differences in the frameworks translate into differences in debt growth reflects very comparable dynamics of long-term interest rates across regimes under this scenario.

5. Robustness analysis in alternative models

Bank of Canada staff are also exploring the robustness of the results presented in models that capture channels that are either absent from ToTEM or present but in limited form. In particular, Djeutem, Reza and Zhang (forthcoming) explore how allowing for a richer treatment of household heterogeneity might alter the findings presented earlier. They do this by using a simple Heterogeneous-Agent New Keynesian (HANK) model following Acharya and Dogra (2020) and Acharya, Challe and Dogra (2020). This model allows them to study the implications of alternative frameworks for variations in inequality over the business cycle. In addition, Wagner, Schlanger and Zhang (forthcoming) study the potential implications of a fuller departure from rational expectations by using a simple New Keynesian model with bounded rationality similar to that of Gabaix (2019, 2020). The analysis in both models is based on an approach that is analogous to the one used in Section 4. We present simulation results in both models without considering the ELB. Results with the ELB will be published in the aforementioned forthcoming papers.

Simple HANK model: brief description and results

We use a simple HANK model as a first step in thinking about the distributional implications of the alternative frameworks. The model features a continuum of households facing uninsurable idiosyncratic income risk. This risk leads to consumption inequality and precautionary saving incentives—two elements that simpler New Keynesian models usually abstract from. Among several important corollaries is that precautionary saving incentives tend to amplify the effects of monetary policy. All else equal, less variation in the interest rate is thus required to achieve the same movements in the output gap and inflation.

Table 4 presents the unconditional standard deviations of several key variables under each framework in the absence of the ELB. As in the case without the ELB in ToTEM, no framework strictly dominates the others. Moreover, PLT remains the framework that best stabilizes inflation, while AIT best stabilizes the output gap. In contrast to the results in ToTEM, NGDP level targeting dominates both FIT and AIT in stabilizing inflation. Moreover, PLT dominates the other frameworks in terms of its ability to economize on changes in the interest rate, which is not the case in ToTEM. The absence of RoT price setters and the fact that the simple HANK model has only a single price level (i.e., no distinction between consumer prices and the GDP deflator) could explain these differences.

Table 4: Unconditional standard deviations (%) in the HANK model					
Policy rule	**Inflation**	**Output gap**	**First difference of interest rate**	**Consumption inequality**	**Relative loss**
Flexible inflation targeting	0.57	0.46	1.28	1.10	1.00
Average inflation targeting	0.50	0.39	1.33	1.16	0.82
Price-level targeting	0.30	0.69	1.23	2.09	0.55
Nominal-GDP level targeting	0.32	0.60	1.27	1.95	0.53
Nominal-GDP growth targeting	0.36	0.60	1.43	1.92	0.59

Importantly, the simple HANK model allows us to identify the regimes that perform best in stabilizing consumption inequality. In this setting, frameworks that are more effective at stabilizing the output gap also tend to do a better job stabilizing consumption inequality.[22] As shown in **Table 4**, FIT and AIT regimes perform best in this dimension, while regimes with more history dependence, such as PLT and NGDP level targeting, exhibit greater volatility of inequality. The delegated loss functions used in modelling the frameworks play a role in generating these results. For example, under PLT, the inclusion of the price level in the delegated loss leads the central bank to strongly stabilize inflation at the cost of allowing greater volatility of the output gap and inequality. Moreover, agents in the HANK model are fully rational. Departures from rational expectations could further increase the volatility of the output gap and inequality.

Table 4 also shows the micro-founded welfare losses of business cycles under each alternative framework, expressed relative to the losses obtained under the FIT framework. These losses are given by the following equation:

$$E[L] = var(\hat{\pi}_t) + 0.016var(x_t) + 0.002var\left(\hat{\Sigma}_t\right) + 0.001var(\hat{y}_t) + 0.002cov(\hat{y}_t, \hat{\Sigma}_t)\,, \qquad (7)$$

where x_t denotes the deviation of output from its flexible price level, $\hat{\Sigma}_t$ represents a measure of consumption inequality in deviation from its steady-state level and $\hat{y}_t$ is the level of output in

[22] This finding is conditional on the assumption that income risk is countercyclical. In this model, cyclical variation in inequality depends on both cyclical variation in the output gap and the real interest rate.

deviation from its steady state. As is common in New Keynesian models with Calvo-style pricing, this loss function assigns an extremely large weight to inflation relative to other variables. Using this metric would suggest that NGDP level targeting is the best of the frameworks in question, followed closely by PLT.

New Keynesian model with bounded rationality: brief description and results

We present analysis in a New Keynesian model with bounded rationality to assess the sensitivity of alternative frameworks to the assumption of rational expectations. Agents in this model are myopic, following Gabaix (2020).[23] For example, relative to a benchmark under rational expectations, households in the model tend to over-discount future changes in income and interest rates when forming expectations and making consumption plans. Similarly, firms over-discount future changes in expected inflation and other key variables when making pricing decisions. As a result, the expectations channel is weaker than that in a model with fully rational expectations.

Table 5 presents the unconditional standard deviations of several key variables under each framework without the ELB. Similar to the results obtained in ToTEM, no framework strictly dominates the others. Moreover, PLT remains the framework that best stabilizes inflation, while AIT is the one that best stabilizes the output gap. Frameworks with greater history dependence, such as PLT and NGDP level targeting, do not perform well in stabilizing the output gap, consistent with the weaker expectations channel in this model.

In contrast to the results from ToTEM, NGDP level targeting now dominates both FIT and AIT in stabilizing inflation. This partly reflects the fact that the model's relatively simple structure implies that CPI and the GDP deflator now coincide.

[23] Gabaix (2020) has offered a comprehensive framework in analyzing a New Keynesian model with bounded rationality. In particular, Gabaix models agents' partial myopia toward distant atypical events using a new micro-founded "cognitive discounting" parameter. The welfare-based loss function in this model is the same as in the canonical New Keynesian model because myopic households continue to experience utility as they would under rational expectations.

Table 5: Unconditional standard deviations (%) in model with bounded rationality				
Policy rule	**Inflation**	**Output gap**	**First difference of interest rate**	**Relative loss**
Flexible inflation targeting	0.45	0.51	0.96	1.00
Average inflation targeting	0.46	0.50	0.99	1.04
Price-level targeting	0.34	0.95	0.87	0.73
Nominal-GDP level targeting	0.39	0.75	1.05	0.81
Nominal-GDP growth targeting	0.37	0.71	1.42	0.73

5. Concluding remarks

The analysis summarized in this document does not identify a clear winner of the horse race, but it does help to identify the relevant trade-offs. FIT, AIT and the unemployment-inflation dual mandate perform well along a number of margins, though none dominates across the board. In contrast, PLT and NGDP level targeting perform well on a narrower set of margins. PLT does a good job of stabilizing inflation, but this comes at the cost of greater output volatility and greater cyclical volatility in inequality. The high degree of history dependence inherent in PLT and NGDP level targeting causes them to do relatively well in ELB episodes but leads to destabilizing behaviour in other situations. In contrast, NGDP growth targeting performs very poorly in all situations.

Overall, FIT, AIT and the dual mandate stand out as the most robust frameworks, performing well for a wide range of different shocks, models and assumptions. Importantly, the differences in performance among frameworks are small by historical standards. Consequently, in addition to studying alternative frameworks, it would be worthwhile to examine the scope for capturing key elements of AIT and the dual mandate within a FIT framework.

References

Acharya, S., E. Challe and K. Dogra. 2020. "Optimal Monetary Policy According to HANK." Federal Reserve Bank of New York Staff Report No. 916.

Acharya, S. and K. Dogra 2020. "Understanding HANK: Insights from a Prank." *Econometrica*, 88 (3): 1113–1158.

Alpanda, S. and S. Zubairy. 2017. "Addressing Household Indebtedness: Monetary, Fiscal or Macroprudential Policy?" *European Economic Review* 92 (C): 43–73.

Amano, R., T. J. Carter and L. L. Schembri. 2020. "Strengthening Inflation Targeting: Review and Renewal Processes in Canada and Other Advanced Jurisdictions." Bank of Canada Staff Discussion Paper No. 2020-7.

Amano, R., S. Gnocchi, S. Leduc and J. Wagner. 2020. "Average Is Good Enough: Average-Inflation Targeting and the ELB." Bank of Canada Staff Working Paper No. 2020-31.

Ambler, S. 2020. "Nominal GDP Level Targeting." Paper prepared for McGill University Max Bell School Monetary Policy Online Conference, "Choosing the Right Target: Real Options for the Bank of Canada's Mandate Renewal," September 22–25.

Andrés, J., D. López-Salido and E. Nelson. 2004. "Tobin's Imperfect Asset Substitution in Optimizing General Equilibrium." The Federal Reserve Bank of St. Louis Working Paper No. 2004-003.

Bank of Canada. 2016. Renewal of the Inflation-Control Target: Background Information. October.

Bank of Canada. 2021. "Toward 2021: Consultations with Canadians." March 31.

Bernanke, B. S. 2017. "Monetary Policy in a New Era." Paper presented at the Peterson Institute for International Economics Conference, "Rethinking Macroeconomic Policy," Washington, DC, October 12–13.

Calvo, G. A. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12 (3): 383–398.

Carter, T. J., R. R. Mendes and L. L. Schembri. 2018. "Credibility, Flexibility and Renewal: The Evolution of Inflation Targeting in Canada." Bank of Canada Staff Discussion Paper No. 2018-18.

Chen, H., V. Curdia and A. Ferrero. 2012. "The Macroeconomics Effects of Large-scale Asset Purchase Programmes." *The Economic Journal* 122 (564): F289–F315.

Christiano, L. J., M. Eichenbaum and C. L. Evans. 2005. "Nominal Rigidities and the Dynamic Effect of a Shock to Monetary Policy." *Journal of Political Economy* 113 (1): 1–45.

Coletti, D., R. Lalonde and D. Muir. 2008. "Inflation Targeting and Price-Level-Path Targeting in the Global Economy Model: Some Open Economy Considerations." *IMF Staff Papers* 55, (2): 326–338.

Corrigan, P., H. Desgagnés, J. Dorich, V. Lepetyuk, W. Miyamoto and Y. Zhang. 2021. "ToTEM III: The Bank of Canada's Main DSGE Model for Projection and Policy Analysis." Bank of Canada Technical Report No. 119.

Djeutem, E., A. Reza and Y. Zhang. Forthcoming. "Household Heterogeneity and the Performance of Monetary Policy Frameworks." Bank of Canada Staff Working Paper.

Gabaix, X. 2019. "Behavioral Inattention." In *Handbook of Behavioral Economics and Applications*, Vol. 2, edited by D. Bernheim, S. DellaVigna and D. Laibson, 261–343. Oxford: Elsevier.

Gabaix, X. 2020. "A Behavioral New Keynesian Model." *American Economic Review* 110 (8): 2271–2327.

Galí, J. and M. Gertler. 1999. "Inflation Dynamics: A Structural Econometric Analysis." *Journal of Monetary Economics* 44 (2): 195–222.

Mendes, R. and S. Murchison. 2014. "Should Forward Guidance Be Backward-Looking?" *Bank of Canada Review* (Autumn).

Smets, F. 2003. ''Monetary Price Stability: How Long Is the Medium Term?'' *Journal of Monetary Economics* 50 (6): 1293–1309.

Smets, F. and R. Wouters. 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review* 97 (3): 586–606.

Svensson, L. E. O. 1999. "Price-Level Targeting versus Inflation Targeting: A Free Lunch?" *Journal of Money, Credit and Banking* 31 (3): 277–295.

Swarbrick, J. and Y. Zhang. Forthcoming. "A Horse Race of Monetary Policy Strategies for Canada." Bank of Canada Staff Discussion Paper.

Wagner, J., T. Schlanger and Y. Zhang. Forthcoming. "A Horse Race of Alternative Monetary Policy Regimes Under Bounded Rationality." Bank of Canada Staff Discussion Paper.

Wilkins, C. A. 2018. "Choosing the Best Monetary Policy Framework for Canada." Speech at McGill University Max Bell School of Public Policy, Montreal, Quebec, November 20.

Woodford, M. 2003. *Interest and Prices: Foundation of a Theory of Monetary Policy*. Princeton: Princeton University Press.