The Bank of Canada's "Horse Race" of Alternative Monetary Policy Frameworks: Some Interim Model Simulation Results[1,2]

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.[3] The agreement has been renewed every few years.[4] The most recent agreement was signed in 2016 and runs to the end of 2021. The agreement defines the inflation target as the 2 percent mid-point of a 1 to 3 percent control range for the 12-month rate of change of the total 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 and the merits of changing the level of the inflation target or adopting price-level targeting.

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

The priorities for the 2021 renewal were established in 2018 with two key challenges in mind (Wilkins, 2018):

(1) The low neutral rate of interest, which reduces the scope to cut the conventional policy interest rate when needed.
(2) Low interest rates associated with a low neutral rate may encourage excessive risk-taking and debt accumulation.

In order to address these challenges, the review includes three streams of work:

(1) Running a "horse race" of alternative monetary policy frameworks (i.e., alternatives to the 2 percent inflation target).

[1] 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.
[2] A broad set of Bank of Canada staff contributed to the work summarized herein, including Robert Amano, Vivian Chu, Paul Corrigan, Edouard Tsague 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. This summary was drafted by José Dorich, Rhys Mendes and Yang Zhang.
[3] See Amano, Carter and Schembri (2020) and Carter, Mendes and Schembri (2018).
[4] 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.

(2) Ensuring that the Bank has an appropriate range of monetary policy tools to achieve its objective(s).
(3) Assessing complementarities between monetary, fiscal and macroprudential policies.

This document summarizes interim results from the first stream — the horse race of alternative frameworks. This stream builds on work from past reviews that looked at changes to the objective of monetary policy. The work for the current review differs from past reviews in its breadth. In addition to revisiting the case for raising the inflation target, 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 employment-inflation dual mandate and nominal GDP (NGDP) growth and level targeting.

Some of these alternatives, such as PLT, imply greater history dependence than conventional inflation targeting. In principle, history dependence can lead to better performance of monetary policy in a low neutral rate environment, but it can lead to greater output volatility if economic agents don't have rational expectations. Other frameworks, such as the employment-inflation dual mandate and NGDP targeting place more explicit emphasis on stabilization of a specific real variable than does inflation targeting. However, inflation targeting as practiced in Canada and elsewhere is flexible inflation targeting. That is, central banks take account of the real side of the economy in pursuing their inflation targets even if no real objective is explicitly identified 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 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 the implications for accountability, communications and credibility. The assessment will involve several different types of work:

- Simulations in several different macroeconomic models
- Laboratory experiments to assess how well real people understand the alternative frameworks
- Empirical work to assess the historical impact of inflation targeting on labour market outcomes
- Public consultations

Much of this work is still ongoing and will be shared in the coming months. This document summarizes the interim results from the simulation analysis in three macroeconomic models. The starting point for this analysis is ToTEM, one of the Bank's main policy analysis and projection models. The ToTEM results are complemented by analysis in two smaller models that focus on specific issues. In particular, we use these models to evaluate the impact of heterogeneity and Gabaix-style bounded rationality.

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

- The performance of the first three frameworks under consideration – IT, AIT and PLT – critically depends on the importance of the ELB constraint.
    - In the absence of the ELB, or if the central bank's extended toolkit can effectively offset the impact of the ELB, 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, IT tends to dominate AIT and PLT in this situation.
    - However, 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, more history dependent policies automatically keep the policy rate at the ELB for a longer period of time. AIT best balances this effect with the additional volatility that history dependence can produce away from the ELB.
- When broadening the horse race to include frameworks that place more explicit emphasis on stabilizing a real variable – the unemployment-inflation dual mandate, NGDP growth targeting and NGDP level targeting – some additional results emerge.
    - Among the additional entrants in the horse race, the unemployment-inflation dual mandate outperforms both variants of NGDP targeting on most dimensions.
    - Among the full set of frameworks, IT, AIT and the unemployment-inflation dual mandate stand out as the most robust of the frameworks in the horse race.
    - Quantitatively, the differences in unconditional volatilities across frameworks are small relative to the shifts seen over history.
    - However, more history-dependent frameworks tend to perform better in large downside scenarios and reduce the frequency of very negative inflation outcomes at the ELB.
- The analysis in the models with heterogeneous agents and bounded rationality highlights potential costs associated with very high degrees of history dependence.
    - In the model with heterogeneous agents, the cyclical variation in inequality is closely related to the cyclical variation in the output gap. The very history-dependent frameworks tend to generate more output gap volatility and therefore lead to greater cyclical variation in inequality.
    - The bounded rationality model reinforces the sensitivity of highly history-dependent frameworks to assumptions about expectations formation.

In ongoing work, we are considering other simulation analysis, which will also play an important role in the ultimate assessment of the alternative frameworks. This analysis includes:

- Evaluating the implications of alternative frameworks for financial stability.
- Revisiting the implications of raising the inflation target in environments in which the long-run Phillips curve is not vertical and there can be hysteresis in the labour market.

- Explicitly modelling the extended monetary policy toolkit.
- Allowing for forward-looking behaviour in the policy rules used to characterize frameworks.
- Modelling frameworks as delegated loss functions that the central bank is tasked to minimize in a discretionary manner.[5]
- Studying the performance of alternative objectives as temporary thresholds.[6]

How well the public understands the alternative frameworks and how private-sector behaviour adjusts will also be important for the overall assessment. The ongoing laboratory experiments and public consultations will be informative about these considerations.

The remainder of this document is organized as follows. Section 1 provides an overview of ToTEM, the benchmark model used in the horse race. Section 2 compares IT, AIT and PLT. Section 3 expands the horse race to include the unemployment-inflation dual mandate, NGDP growth targeting and NGDP level targeting. It also provides some robustness analysis from other models. We offer some concluding remarks in Section 4.

1. Brief description of ToTEM

ToTEM is a large-scale open economy DSGE model of the Canadian economy. 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 GDP, which is important for an accurate evaluation of monetary policy frameworks targeting the level or growth rate of nominal GDP or incorporating some role for the output gap.

The firms responsible for producing final goods face nominal rigidities when setting their prices. More specifically, some of the firms in a given final good producing sector reoptimize 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 similar to that in Galí and Gertler (1999). Sector-specific shares for each of these two pricing types have been estimated. The estimated share of rule-of-thumb price setters is relatively high in some sectors. For example, the share of rule-of-thumb price-setters in the core consumption sector is estimated to be about 50 percent. This is especially important because higher shares are well known to undermine the performance of price-level targeting and other history-dependent regimes that rely more heavily on expectational channels.

[5] This alternative approach avoids constraining the reaction function to take the form of a simple rule.
[6] See Bernanke (2017) and Mendes and Murchison (2014).

Turning to the household block of ToTEM, we note that the model features three prominent household types differing in terms of the financial markets to which they have access and their status as borrowers or savers in those markets. On the saver side, the model follows Andres, Lopez-Salido and Nelson (2004) and Chen, Curdia and Ferrero (2012) in assuming that some savers are "restricted" in the sense that they cannot access long-term debt markets, while others are "unrestricted" with access to all 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, a feature of the model that could be particularly important when evaluating in future research the complementarities between large scale asset purchases and the frameworks considered here. Taken together, the two saver types 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 "unrestricted" saver households. When doing so, 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, while the second aims to capture home equity lines of credit. Since these two funding sources account for more than 80 per cent of total household debt in Canada, the model is able to offer 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 a rule of thumb. Shares of the two types have been estimated, and the sizeable share of rule-of-thumb types that has been estimated points to this feature of the model as another quantitatively important dimension in which ToTEM departs from the textbook assumption of fully rational, forward-looking behaviour. As mentioned 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. A simple equation relating unemployment to hours worked has therefore been estimated and added 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

[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. This type accounts for less than five percent of all households.

as movements in labour force participation and the efficiency of the matching process between job-seekers and vacancies.

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.

2. ToTEM Results (Part I): IT, AIT and PLT

The first three frameworks we consider differ only in the degree of history dependence embedded in the targeted price variable. Inflation targeting, average inflation targeting and price-level targeting all involve targeting a variable of the form:

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

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 inflation targeting 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, price-level targeting can be viewed as a limiting case. Intermediate cases ($12 < N < \infty$) correspond to average inflation targeting – 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 very history-dependent target variable 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.

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

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

where i_t is the nominal policy interest rate, $\tilde{x}_t$ is the output gap, $\bar{\pi}$ is the inflation target expressed at a monthly rate under IT and AIT and the trend inflation rate under PLT.[8]

For concreteness, we consider three cases:

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

[8] We have in mind PLT regimes in which there can be a deterministic trend in the targeted path for the price level.

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

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

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, $\bar{\pi}^a$ is a constant target set at 2 percent in our simulations. The price level target has a deterministic trend: $\bar{p}_t = \bar{p}_{t-1} + \bar{\pi}$. 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. In addition, we had studied inflation targets above 2 percent for the 2016 renewal (Bank of Canada, 2016) and we are looking at that issue again in ongoing work. Among other things, this work includes using a model with a non-vertical long-run Phillips curve.

The benefits of history dependence

In order 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 IT 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 in order to force inflation below trend to return the price-level to its target. Given these results, the choice between IT and PLT reduced to a judgment call on 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. 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, the nature of the regime must be understood by economic agents 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 are *eventually* 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 IT, AIT and PLT are all based on consumer prices in the model. The Bank of Canada's actual IT 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 (3)-(5) in order to capture the idea that the Bank can look through temporary volatility.[9] Note also that there is a unique concept of core in the model, so we do not have to choose among the Bank's three empirical measures of core inflation.

We consider the performance of the rules in equations (3)-(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 make a clean comparison of the 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) toolkit (QE, etc.) at the ELB are important for our evaluation of alternative frameworks. In work for the 2016 renewal, Bank staff had explicitly modelled the effects of the EMP tools. That approach is being implemented in ongoing work but is not reported in this document. Rather, we deal with the EMP toolkit 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 extended monetary policy toolkit 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).[10] This can be viewed as approximating a situation in which the EMP toolkit 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 toolkit. Reality lies somewhere in between, with the EMP toolkit only able to partially substitute for

[9] 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.
[10] We have in mind the concept of a shadow rate in the sense of Black (1995) as applied by Wu and Xia (2016).

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 IT, AIT and PLT

We begin by assuming that the central bank aims to minimize the variances of 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. **Figure 1** plots frontiers both with and without the ELB. These frontiers are derived by searching over the γ and α parameters and plotting the lowest achievable pairs of the unconditional variances of the output gap and inflation. The frontiers suggest that:

- In the absence of the ELB (solid lines in **Figure 1**), standard inflation targeting performs very well. Given our interpretation of the "No ELB" case, this suggests that IT would do well in a situation in which the EMP toolkit 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 **Figure 1**), IT is always dominated by AIT and it is dominated by PLT for parameterizations that lead to low inflation volatility. The weaker performance of IT in this case reflects the fact that history dependence is more important when the ELB binds. In this situation, more history dependent policies automatically keep the policy rate at the ELB for a longer period of time. This provides additional stimulus that mitigates the effect of the ELB.

Figure 1: Policy Frontiers

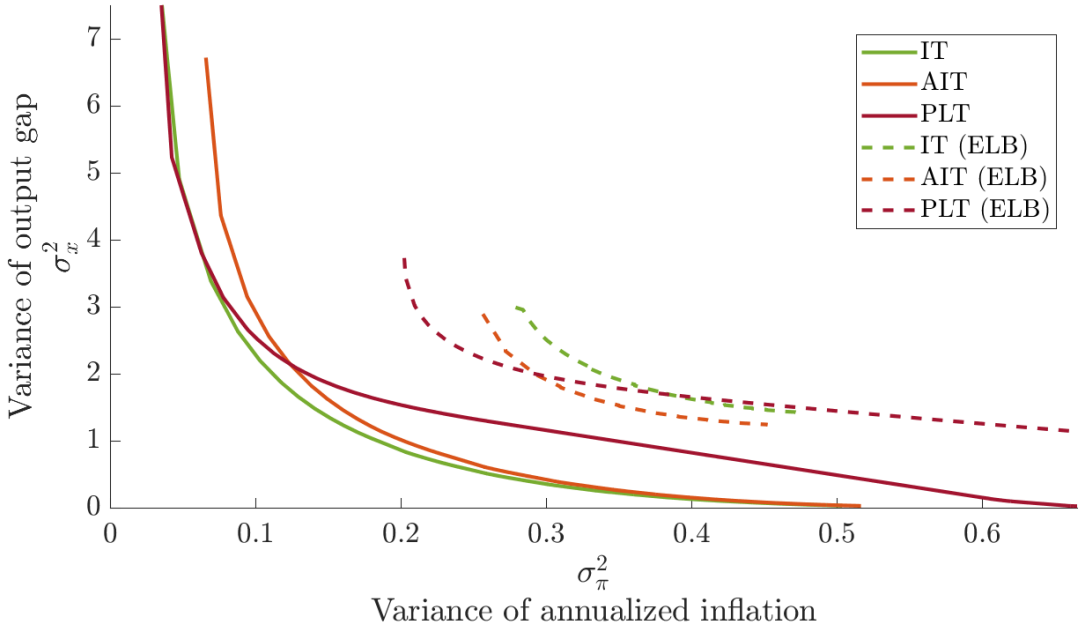

These results from ToTEM contrast with the dominance of PLT in the canonical New Keynesian model. The difference stems from the high proportion of rule-of-thumb wage- and price-setters in ToTEM. This is qualitatively consistent with the findings of Amano et al. (2020). They show that, in the presence of rule-of-thumb firms, highly history-dependent frameworks generate a worse inflation-output trade-off. Thus, rule-of-thumb price-setters reduce the optimal degree of history dependence.[11] Amano et al. 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 **Figure 1**.

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

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

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 interest rate as before, $\bar{\pi}^a$ is the inflation target and λ_x is the relative weight on the 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 λ_x we select the γ and α parameters in equations (3)-(5) to minimize the loss function in (6). **Figures 2** and **3** plot the value of the loss function for each framework for values of λ_x between 0 and 2.

Figure 2 focuses on the case in which the ELB is not a constraint (or the EMP toolkit is very effective). In this case, IT generates the lowest loss for $\lambda_x > 0.26$. This suggests that, if the EMP toolkit is effective, an inflation target would perform best in an environment where the central bank places significant emphasis on stabilizing the real economy. The more history dependent policies (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 rule-of-thumb 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 both inflation and output gap variability. That said, the differences between IT and AIT are quantitatively small regardless of the value of λ_x.

[11] Our results assume that the proportion of rule of thumb 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 of Canada staff are looking at this question in a model that allows firms to choose whether or not to follow a simple rule of thumb.

Figure 2: Loss Values with No ELB

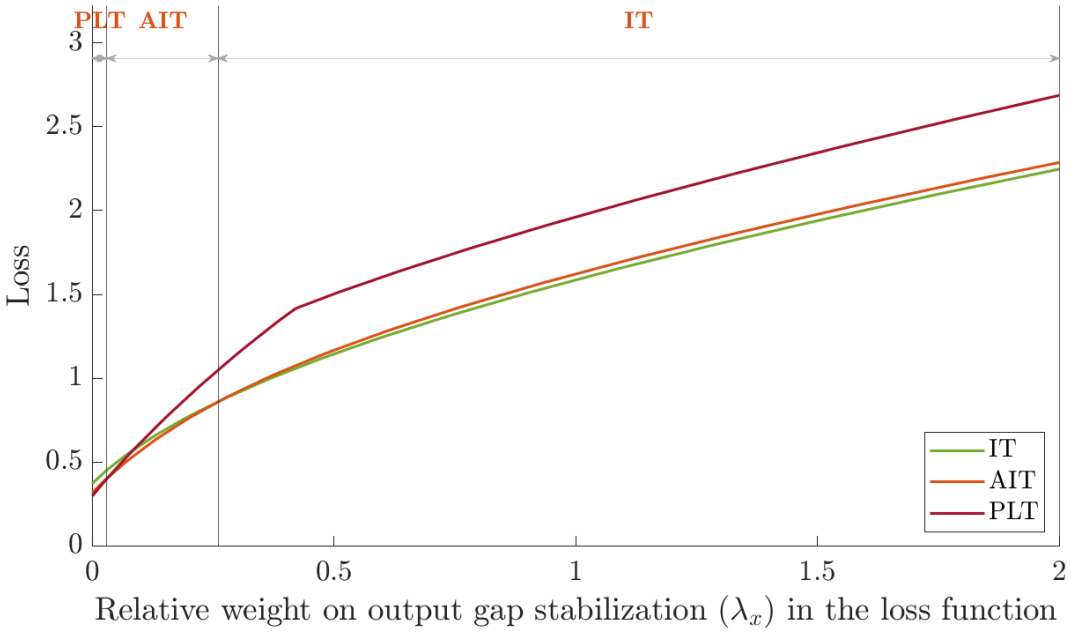

Figure 3 shows the losses in the case with an occasionally binding ELB. With the ELB, history dependence yields greater benefits. IT is now dominated by AIT and PLT for all values of λ_x. AIT yields the lowest loss for a wide range of λ_x values. In other words, if the EMP toolkit is ineffective, an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

Overall, the results in this section suggest that PLT is dominated by less history-dependent alternatives regardless of how much relative weight the central bank places on stabilizing the real economy.[12] The "horse race" between IT and AIT depends on the assumed effectiveness of the EMP toolkit. If EMP tools are effective at compensating for the ELB, our results suggest that there are no gains from moving from IT to AIT. However, even when EMP tools are effective, IT is only marginally superior to AIT. On the other hand, AIT has a somewhat larger lead over IT if EMP tools are assumed to be ineffective. Thus, AIT might be regarded as more robust to the effectiveness of the EMP toolkit given that its relative performance is good in both cases.

[12] 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).

Figure 3: Loss Values with Occasionally Binding ELB

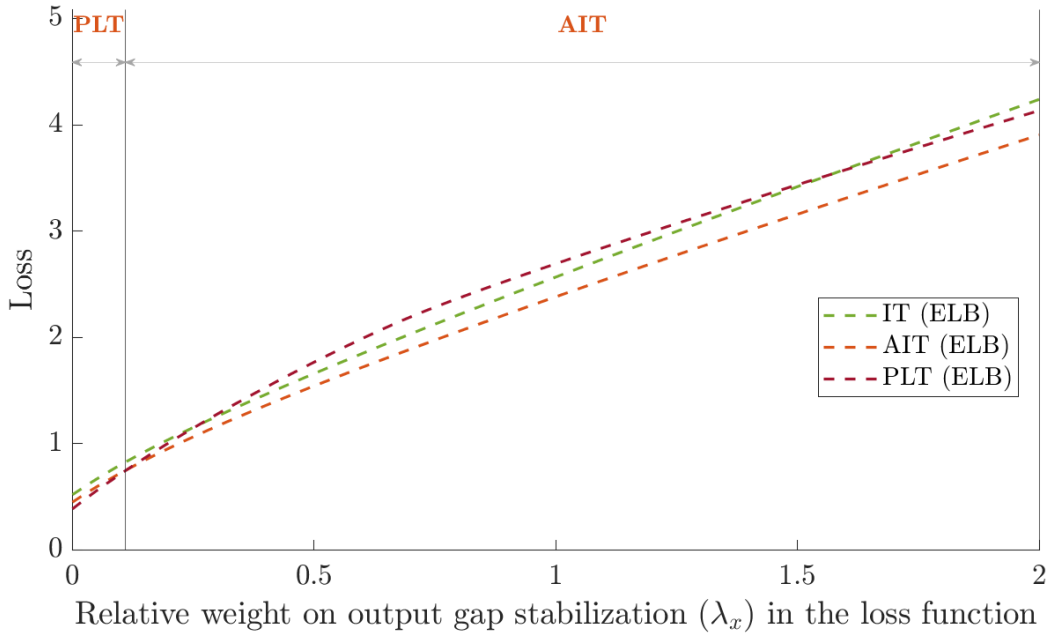

3. ToTEM Results (Part II): Expanding the Horse Race

The Expanded 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 that we need a different approach to run a fair horse race with this broader set of frameworks.

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

(1) There are significant technical obstacles to 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 introduce microfoundations that have unrealistic implications for welfare.

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

- Each framework is characterized by a regime-specific loss function delegated to the central bank and a simple policy rule (**Table 1**). The delegated loss function is used to choose the parameters of the rule, but it is not used to evaluate the framework.
- The frameworks are evaluated using volatilities of several key economic variables. We do not explicitly weight these volatilities into a loss function. Rather, we 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[13]
IT	$L^{IT} = \left(\pi_t^{yy} - \bar{\pi}^a\right)^2 + 0.5\,(\Delta i_t)^2 + (\widetilde{x_t})^2$	$i_t = i^* + \gamma\left(\pi_t^{yy} - \bar{\pi}^a\right) + \alpha\widetilde{x_t}$
AIT	$L^{AIT} = \left(\pi_t^{3y} - \bar{\pi}^a\right)^2 + 0.5\,(\Delta i_t)^2 + (\widetilde{x_t})^2$	$i_t = i^* + \gamma\left(\pi_t^{3y} - \bar{\pi}^a\right) + \alpha\widetilde{x_t}$
PLT	$L^{PLT} = (p_t - \bar{p}_t)^2 + 0.5\,(\Delta i_t)^2 + (\widetilde{x_t})^2$	$i_t = i^* + \gamma(p_t - \bar{p}_t) + \alpha\widetilde{x_t}$
NGDP Level	$L^{NGDPL} = \begin{Bmatrix} (y_t + p_{GDP,t}) \\ -(\bar{y}_t + \overline{p_{GDP,t}}) \end{Bmatrix}^2 + 0.5\,(\Delta i_t)^2$	$i_t = i^* + \delta\begin{bmatrix} (y_t + p_{GDP,t}) \\ -(\bar{y}_t + \overline{p_{GDP,t}}) \end{bmatrix}$
NGDP Growth	$L^{NGDPG} = \begin{Bmatrix} (\Delta y_t^{yy} + \Delta p_{GDP,t}^{yy}) \\ -(\overline{\Delta y_t}^{yy} + \overline{\Delta p_{GDP,t}}^{yy}) \end{Bmatrix}^2 + 0.5(\Delta i_t)^2$	$i_t = i^* + \delta\begin{bmatrix} (\Delta y_t^{yy} + \Delta p_{GDP,t}^{yy}) \\ -(\overline{\Delta y_t}^{yy} + \overline{\Delta p_{GDP,t}}^{yy}) \end{bmatrix}$
Unemployment-Inflation Dual Mandate	$L^{DM} = \left(\pi_t^{yy} - \bar{\pi}\right)^2 + (\widetilde{u_t})^2 + 0.5\,(\Delta i_t)^2$	$i_t = i^* + \gamma\left(\pi_t^{yy} - \bar{\pi}\right) + \alpha\widetilde{u_t}$

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 US Federal Reserve 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 IT, except that the output gap measure in the loss function and interest rate rule is replaced by an unemployment gap, $\widetilde{u_t}$, measuring the difference between the actual and natural rates of unemployment.

[13] In the simulations, all the frameworks considered in Table 1 assume a common smoothing parameter of 0.85, similar to equations (3)-(5).

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, $p_{GDP,t}$. This framework has received renewed attention because of its potential to address some of the challenges of the current environment.[14] 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 actual and potential 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 IT, 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 targeting has important disadvantages too. For example, if changes in trend real growth are not offset by changes in the target path for nominal GDP, then trend inflation will be forced to adjust. This would reduce agents' level of certainty about future inflation when making long-term decisions and it 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 targeting.

Nominal GDP Growth Targeting: We model NGDP growth targeting as stabilizing the sum of the logarithm of year-over-year real GDP growth, Δy_t^{yy}, and the logarithm of the year-over-year rate of change of the GDP deflator, $\Delta p_{GDP,t}^{yy}$. This framework differs from IT in two important dimensions. First, it incorporates the rate of change of GDP deflator, rather than consumer price 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 differing in their treatments of the effective lower bound (ELB) on nominal interest rates and the nature of the extended toolkit assumed available to the central bank. As in the previous section, our first case accounts for the ELB but assumes that the extended toolkit 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 toolkit can perfectly substitute for conventional monetary stimulus, thus allowing the central bank to achieve the same outcomes as in the absence of the ELB. For this reason, we interpret negative rates in this case as representing a shadow rate. Together, these two cases represent natural benchmarks in the sense that they provide upper and lower bounds on a given framework's performance in a more realistic scenario where (i) the central bank has access to an extended toolkit of some form, but (ii) the instruments therein cannot fully offset the impact of the ELB. A more explicit analysis accounting for the

[14] For instance, see Ambler (2020).

likely availability and effectiveness of these instruments in the context of the horse race will be provided in future work.

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 (Y/Y), the output gap, real GDP growth (Y/Y), the unemployment rate, the first difference of the nominal interest rate and real household debt growth. 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, nominal GDP growth 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 only be favoured under a ranking system that places nearly singular weight on the stabilization of household debt growth.

Table 2: Unconditional standard deviations (%), with ELB

Policy Rule	Total CPI inflation	Output gap	First difference of interest rate	Unemployment rate	Real GDP growth (Y/Y)	Real household debt growth
IT	0.67	1.25	0.69	0.89	1.94	5.69
AIT	0.68	1.20	0.71	0.87	1.94	5.69
DM	0.70	1.42	0.57	0.78	2.07	5.68
PLT	0.56	1.56	0.78	0.99	2.19	5.83
NGDP level	0.84	1.58	1.21	0.99	2.01	5.68
NGDP growth	1.18	2.27	1.22	1.34	2.10	5.50

Similarly, price level targeting emerges as the framework that best stabilizes CPI inflation. Moreover, it is the regime that minimizes the likelihood of very low inflation outcomes (**Figure 4**). However, it is strictly dominated by IT, AIT and the dual mandate with respect to all other variables in the table. This suggests that price level targeting would only be favoured under a ranking system that places very high weight on inflation stabilization. At the same time, IT, AIT and the dual mandate all outperform nominal GDP 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 IT, AIT and the dual mandate.

IT, AIT and the dual mandate thus generally emerge as the most robust of the frameworks in question. A key corollary is that an intermediate degree of history dependence best strikes a reasonable balance between (i) the benefits that history dependence has to offer in terms of its stabilizing effects on expectations at the ELB, and (ii) the costs stemming from its destabilizing interactions with rule-of-thumb agents. That said, we note the quantitative differences in unconditional volatilities across the frameworks under consideration are small relative to historical shifts. In particular, there were substantial declines in macroeconomic volatility after

the adoption of inflation targeting in Canada. For example, the standard deviation of inflation fell from 3.1 percent in the 1970s and 1980s to 0.9 percent in the 1995 to 2016 period.[15] Similarly, the standard deviation of real output growth declined from 3.9 percent to 2.5 percent over that period. In contrast, most of the differences in **Table** 2 are much smaller.

Figure 4: Distribution of Total Inflation under Different Policy Regimes

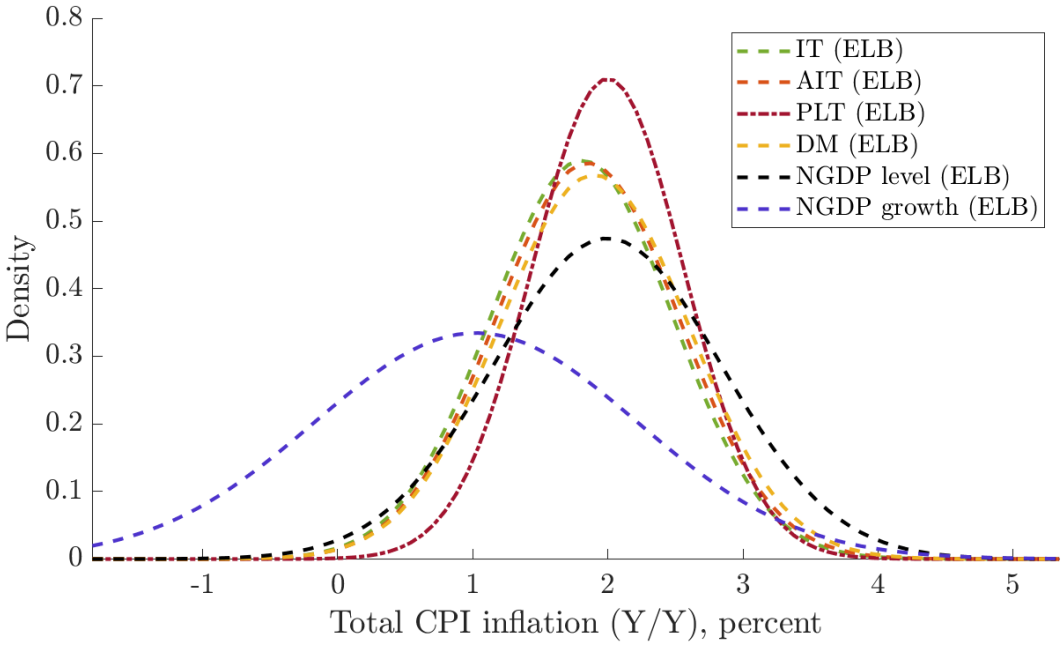

Table 3 presents our results for the case in which we abstract from the ELB. As in the previous case, no one framework dominates the others with respect to all variables under consideration. Moreover, many of the general patterns from our previous case continue to hold in some form. For example, though nominal GDP growth targeting now emerges as the framework that best stabilizes *both* real GDP growth and real household debt growth, its advantages relative to IT 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 Nominal GDP growth targeting some distance behind IT 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 Nominal GDP 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 rule-of-thumb behaviour in the economy and the limits that this behaviour places on the strength of expectational transmission mechanisms.

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

Table 3: Unconditional standard deviations (%), without ELB						
Policy Rule	**Total CPI inflation**	**Output gap**	**First difference of interest rate**	**Unemployment rate**	**Real GDP growth (Y/Y)**	**Real household debt growth**
IT	0.68	0.88	0.87	0.78	1.73	5.64
AIT	0.60	0.79	0.79	0.78	1.64	5.55
DM	0.65	1.30	0.63	0.67	2.01	5.66
PLT	0.51	1.51	0.83	0.97	2.10	5.83
NGDP level	0.81	1.55	1.61	0.97	2.01	5.77
NGDP growth	0.75	1.40	2.17	0.91	1.67	5.59

For these reasons, IT, 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, making the choice among these regimes principally dependent 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, IT 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 toolkit 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 toolkit.

Robustness

In this subsection, we explore the robustness of the results presented above by exploring the various regimes' performance in the context of models that capture channels that are either absent from ToTEM or present but in limited form. In particular, we first explore how allowing for a significantly richer degree of household heterogeneity might alter the findings presented earlier, namely using a simple Heterogeneous-Agent New Keynesian (HANK) model following Acharya and Dogra (2020) and Acharya, Challe and Dogra (2020). We then turn our attention to the potential implications of a fuller departure from rational expectations – an issue that we explore using a simple New Keynesian model with bounded rationality similar to that of Gabaix (2019, 2020).

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 countercyclical, uninsurable idiosyncratic income risk. This risk leads to consumption inequality and precautionary saving incentives, both of which simpler New Keynesian models 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, there is no framework that 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 IT 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 a rule-of-thumb price setters and the fact that there is only a single price level (i.e., no distinction between consumer prices and the GDP deflator) in the simple HANK model could explain these differences.

Table 4: Unconditional standard deviations (%) in HANK model

Policy Rule	Inflation	Output gap	First difference of interest rate	Consumption inequality	Welfare-Based Loss	Relative Loss
IT	0.57	0.45	1.28	1.09	0.094	1.00
AIT	0.50	0.38	1.33	1.15	0.077	0.82
PLT	0.30	0.68	1.23	2.08	0.052	0.56
NGDP level	0.32	0.60	1.28	1.94	0.051	0.54
NGDP growth	0.36	0.60	1.43	1.92	0.056	0.59

Since the simple HANK model includes consumption inequality, we can also use it to identify the regimes that perform best in stabilizing this additional variable. As shown in **Table 4**, IT and AIT are the regimes that perform best in this dimension, while regimes with more history dependence such as PLT and NGDPL exhibit a poorer performance in this regard. Frameworks that are more effective at stabilizing the output gap also tend to do a better job of stabilizing the extent of consumption inequality.

Table 4 also shows the microfounded welfare losses of business cycles under each framework. These losses are given by the following expression:

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

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, $\hat{y}_t$ is the level of output in deviation from its steady state. Interestingly, this loss function assigns an extremely large weight to inflation relative to that on the output gap and consumption inequality. As a result, 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 use a model in which agents are myopic towards future information in order to assess the sensitivity of alternative frameworks to the assumption of rational expectations. In the model, in forming expectations and making consumption plans, households tend to over-discount future changes in income and interest rates relative to the fully rational benchmark on which textbook models focus. Similarly, firms over-discount future changes in expected inflation and other key variables when making pricing decisions. As a result, the expectation 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 in the absence of the ELB. Similar to the results obtained in ToTEM, there is no framework that 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 expectational channel characterizing this model.

In contrast to the results from ToTEM, NGDP level targeting now dominates both IT 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.

Table 5: Unconditional standard deviations (%) in Model with Bounded Rationality

Policy Rule	Inflation	Output gap	First difference of interest rate	Welfare-Based Loss[16]	Relative Loss
IT	0.45	0.51	0.96	0.013	1.00
AIT	0.46	0.5	1	0.014	1.04
PLT	0.34	0.95	0.87	0.01	0.73
NGDP level	0.39	0.75	1.05	0.011	0.81
NGDP growth	0.37	0.71	1.42	0.01	0.73

[16] The welfare characterization of the model with bounded rationality follows that of a canonical New Keynesian model.

Scenario Analysis in ToTEM

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

Figure 5 reports 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 toolkit could be deployed once the nominal rate reaches the ELB.

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, as compared against 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 four quarters at the ELB, while the other regimes imply durations in the range of 9 quarters to 13 quarters. This difference occurs because nominal GDP growth begins to recover about four quarters after the beginning of the simulation, while the level of GDP remains depressed.

As for the other regimes' performance, NGDP level targeting seems to best insulate the economy from the shocks in question, in the sense that the trough responses of CPI inflation, the output gap and real GDP growth are least negative under this regime, while the peak response of the unemployment gap is least positive. This is explained by NGDP level targeting's relatively more aggressive interest rate response, 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.

Figure 5: A Scenario Analysis of Policy Regime Performance in ToTEM

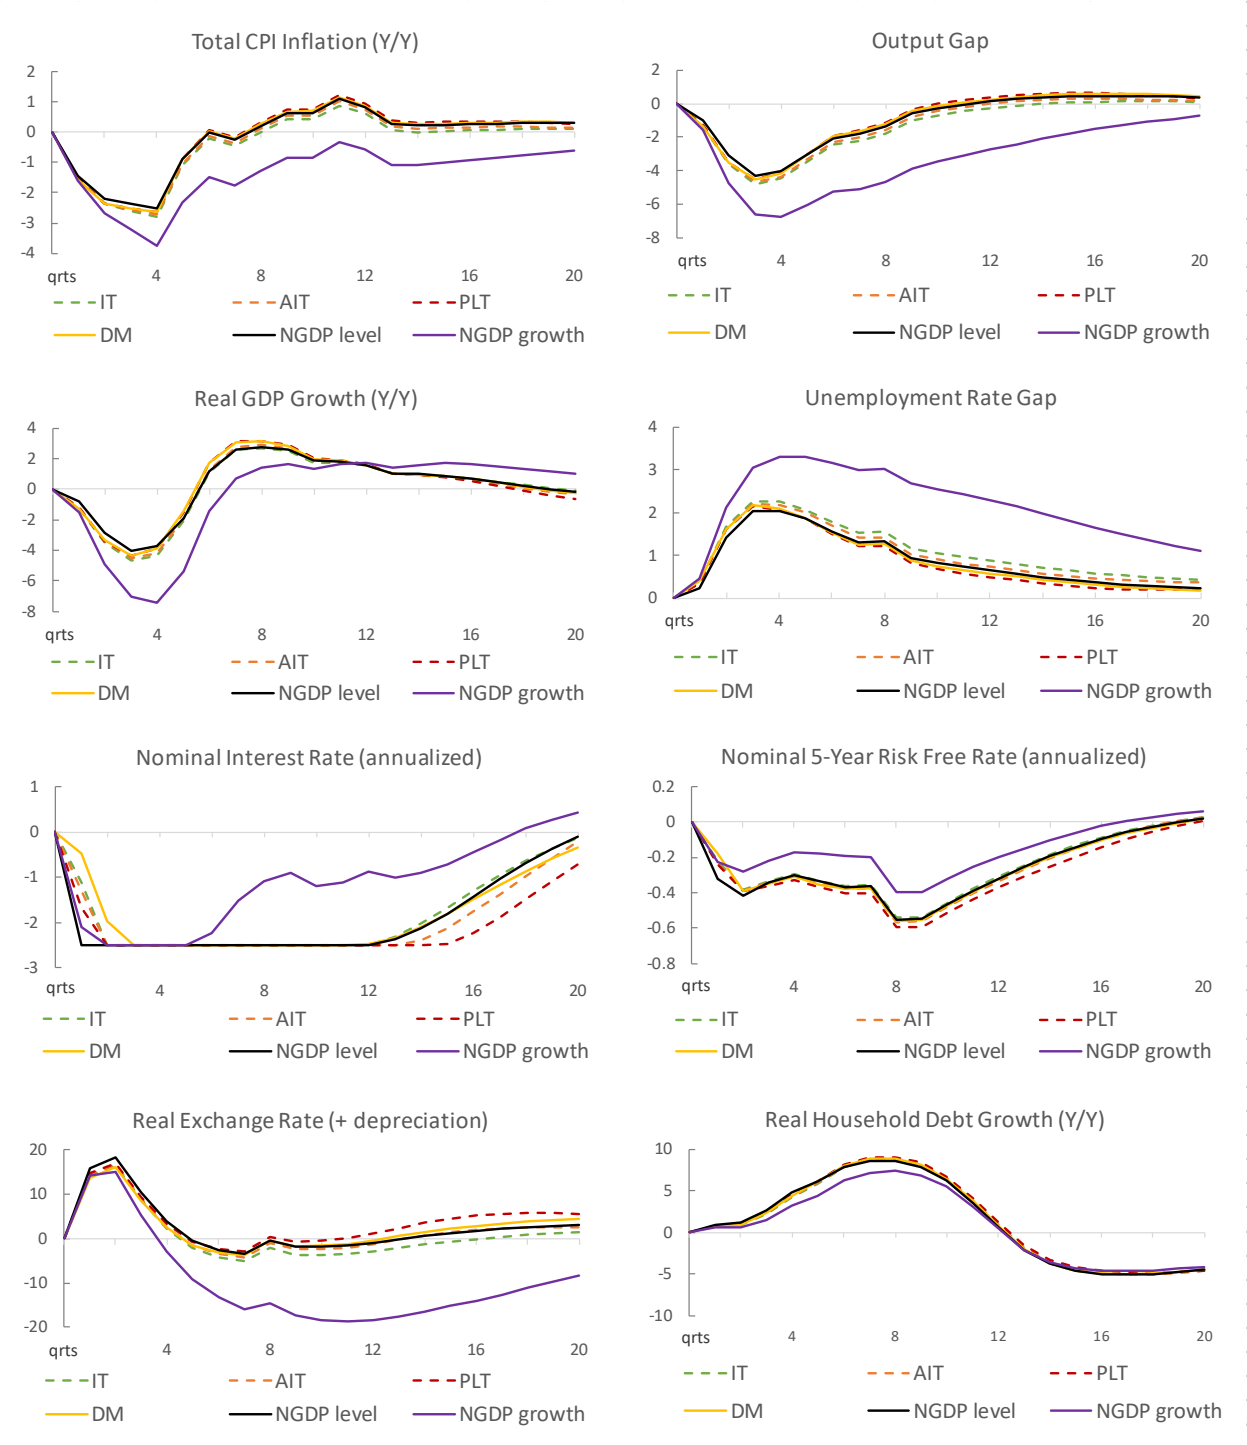

Turning finally to the regimes' potential implications for financial stability, we note that the peak responses of real household debt growth fall within a relatively narrow range of 8.7 to 9.1 percentage points excluding NGDP growth targeting. The relatively limited extent to which differences in the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

4. 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. PLT performs well in terms of inflation stabilization, but this comes at the cost of greater output volatility and therefore greater cyclical volatility in inequality. While the performance of PLT and NGDP level targeting does not stand out on average, both frameworks do well in the case of large negative demand shocks. NGDP growth targeting does not perform well in terms of macroeconomic stabilization. IT, AIT and the unemployment-inflation dual mandate perform well along a number of margins, though none dominate across the board.

Importantly, the interim results suggest that the differences in performance among frameworks are small by historical standards. Consequently, it is not yet clear if any of the alternatives would offer expected gains large enough to justify shifting away from the proven and successful inflation-targeting framework.

The analysis reported in this document focused mainly on the macroeconomic performance of the alternative frameworks. Other modelling work will help to shed additional light on the financial stability and distributional implications of the regimes. We are conducting empirical work to better understand how the inflation targeting framework has affected labour market outcomes, with the aim of informing the analysis of alternative regimes. Through laboratory experiments, we will also gain insights into the comprehensibility of the different frameworks and the implications for expectation formation. Public consultations will help identify what properties of a monetary policy framework are most valued by Canadians. This additional work will help to refine the assessment of the expected benefits and risks associated with each framework.

References

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

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

Alpanda, S., and Zubairy, S. 2017. "Addressing household indebtedness: Monetary, fiscal or macroprudential policy?" European Economic Review, Elsevier, Vol 92(C), pages 43-73.

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

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

Andres, J., Lopez-Salido, D., and Nelson, E. 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.

Bernanke, B.S. 2017. "Monetary Policy in a New Era." Presented at the Peterson Institute for International Economics conference on Rethinking Macroeconomic Policy, October 12-13.

Calvo, G.A. 1983. "Staggered prices in a utility-maximizing framework." *Journal of Monetary Economics,* Elsevier, Vol 12(3), pages 383-398.

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

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

Christiano, L. J., Eichenbaum, M., and Evans, C. L. 2005. "Nominal Rigidities and the Dynamic Effect of a Shock to Monetary Policy." *Journal of Political Economy*, University of Chicago Press, Vol. 113(1), pages 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*, Vol. 55, No. 2.

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

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

Gali, J., and Gertler, M. 1999. "Inflation Dynamics: A Structural Econometric Analysis." *Journal of Monetary Economics*, Elsevier, Vol 44(2), pages 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*, Vol. 50, No. 6, pp. 1293–309.

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

Svensson, L. E. O. 1999. "Inflation targeting as a monetary policy rule." *The Journal of Monetary Economics*, Elsevier, Vol 43(3), pages 607-654.

Svensson, L. E. O. 2020. "Monetary Policy Strategies for the Federal Reserve." NBER Working Paper no 26657.

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 University Press.