
DOI: 10.1111/jmcb.70083

VADYM LEPETYUK
LILIA MALIAR
SERGUEI MALIAR
JOHN B. TAYLOR

The Power of Open-Mouth Policies

Economic agents adjust their behavior in anticipation of announced future policy changes. We develop a novel perturbation-based extended function path framework to construct nonstationary solutions—time-dependent decision functions—for economic models with anticipated, non-Markov news shocks. We apply this framework to evaluate the anticipation effects of policy announcements within a realistically calibrated central banking model. Specifically, we analyze temporary and permanent anticipated changes to policy rules, including shifts in inflation targets and Taylor rule aggressiveness. We also examine anticipated transitions from inflation targeting to price-level and average-inflation targeting, finding that these announcements yield economically significant effects.

JEL codes: C61, C63, C68, E31, E52

Keywords: news shocks, turnpike theorem, time-dependent models, nonstationary models, unbalanced growth, regime switches, monetary policies, price-level targeting, average inflation targeting

CENTRAL BANKS INCREASINGLY RELY ON communication to the public as a key tool for implementing monetary policy. Through these communi-

Lilia Maliar and Serguei Maliar acknowledge financial support from the NSF grants SES-1949413 and SES-1949430, respectively. 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 of Canada.

VADYM LEPETYUK is a Chief Expert, Monetary Policy and Economic Analysis Department, National Bank of Ukraine (E-mail: vadym.lepetyuk@bank.gov.ua). LILIA MALIAR is a Professor of Economics, Economics Department, The Graduate Center, City University of New York (E-mail: lmaliar@gc.cuny.edu). SERGUEI MALIAR is an Associate professor, Economics Department, Santa Clara University (E-mail: maliars@stanford.edu). JOHN B. TAYLOR is a Robert Raymond Professor of Economics, Department of Economics and Hoover Institution, Stanford University (E-mail: john-tayl@stanford.edu).

Received June 15, 2021; and accepted in revised form May 17, 2026.

Journal of Money, Credit and Banking, Vol. 00, No. 0 (August 2026)

© 2026 The Ohio State University.

All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies.

cations, monetary authorities convey their future intentions and provide insights into their views on the economy's future trajectory. For instance, a central bank might commit to maintaining a fixed interest rate for a specific period before normalizing its policy (forward guidance) or announce an upcoming adjustment to the inflation target. Understanding the impact of such communication is crucial for evaluating the effectiveness of monetary policy.

In this paper, we evaluate the impact of central bank announcements by examining several monetary policies proposed in the literature as potential avenues for improving welfare. The policies under consideration include one-time and gradual anticipated changes, as well as more complex scenarios where different monetary policies are implemented with specific probabilities. A key feature of our model, which incorporates communication, is that economic agents respond to policy announcements even when the actual policy changes are scheduled to take effect at a future date. Our findings reveal that these anticipation effects can be substantial.

Our analysis is conducted using a realistically calibrated, prototypical central banking model, which serves as a scaled-down version of the Terms of Trade Model (ToTEM) employed by the Bank of Canada for projection and policy analysis. Notably, our “baby” ToTEM (bToTEM) closely mirrors the structure and implications of the full-sized ToTEM model producing essentially similar impulse response functions to the key shocks. See Lepetyuk, Maliar, and Maliar (2020, henceforth LMM) for detailed description of bToTEM model and Corrigan et al. (2021) for the latest description of ToTEM model.

The policy experiments we examine include the following: (i) a gradual adjustment of the inflation target, scheduled to occur in the future either with certainty or with some probability; (ii) the normalization of monetary policy with respect to future nominal interest rates, when the economy is initially at a zero lower bound (ZLB) on nominal interest rates; (iii) a switch to a more aggressive Taylor rule; (iv) a shift from inflation targeting to price-level targeting (PLT); and (v) a move from inflation targeting to average inflation targeting (AIT). Our findings are summarized as follows:

1. *Inflation target changes*: An increase in the inflation target has a positive short-term effect on output, and delaying this increase by 1 year further amplifies the output gain during the transition to the new steady state. Even when the announced policy is implemented with some probability, substantial anticipation effects persist both before and after the uncertainty is resolved.
2. *Forward guidance*: When the economy is constrained by the ZLB on nominal interest rates, the central bank employs forward guidance—policy announcements about its future return to the standard interest rate rule—to guide the economy out of the ZLB. Extending the timeline for this return results in greater output expansion during the transition. However, the initial jump in output remains unaffected by the length of the forward guidance period. This finding provides policymakers with valuable insights into the optimal duration for monetary policy normalization following ZLB periods.

3. *Aggressive Taylor rule*: A more aggressive (yet realistic) behavior of the central bank in targeting inflation generates limited anticipation effects among economic agents.
4. *Price-level targeting*: Transitioning from inflation targeting to PLT produces smaller impacts when there are longer implementation lags. PLT has been shown in the literature to improve welfare, and thus, delaying its adoption deprives the economy of earlier benefits from higher output.
5. *Average inflation targeting*: Shifting to AIT yields modest anticipation effects, as it represents a middle ground between inflation targeting and PLT.

In summary, our analysis reveals that the significance of anticipation effects varies depending on the policy experiment. Substantial anticipation effects are evident in scenarios involving changes in the inflation target (Experiment 1) and forward guidance at the ZLB (Experiment 2). However, these effects are relatively muted in experiments involving more aggressive monetary policy rules or targeting regimes (Experiments 3–5).

On the methodological side, we argue that announcements about future economic conditions and policies can be effectively modeled as anticipated non-Markov news shocks. These shocks lead to nonstationary solutions (time-dependent optimal decision rules) that cannot be addressed using conventional numerical methods designed for constructing stationary solutions (time-invariant decision rules). Within the framework of perturbation analysis, we modify the traditional approach to enable the construction of time-dependent decision functions. We call our approach the extended function path (EFP), drawing a parallel to the extended path (EP) method introduced by Fair and Taylor (1983). Unlike EP, which constructs paths for variables (time series), EFP constructs paths for decision functions. For further discussion and a review of related literature, see Maliar et al. (2020, henceforth MMTT).

The EFP methodology is applicable to a broad and empirically relevant class of economic models in which finite-horizon trajectories converge to infinite-horizon trajectories as the time horizon extends—a property known as the turnpike theorem. Our perturbation-based EFP method offers accuracy comparable to the global projection method developed in MMTT (2020) while being significantly more computationally efficient for high-dimensional problems, such as large-scale central banking models.

The implementation of our approach is user-friendly, as our MATLAB codes are built upon the widely used Dynare software. This design ensures accessibility and flexibility, allowing users to adapt the method for a variety of applications beyond the scope of this study.

Other methods in the literature allow for the analysis of changes in the economic environment, but these typically rely on the assumption of Markov shocks. For instance, models with Markov regime switching (e.g., Davig and Leeper 2007) and models with Markov news shocks (e.g., Schmitt-Grohé and Uribe 2012) are prominent approaches. A key distinction between these methods and our EFP analysis lies in the treatment of shocks. The Markov-based literature assumes shocks are recurrent, random draws from a stationary Markov distribution, while we focus on

shocks that represent a sequence of events occurring at predetermined dates. As a result, the Markov literature approximates a single, time-invariant decision rule corresponding to the given Markov process. In contrast, our EFP method constructs an infinite sequence of optimal decision rules tailored to a specific sequence of historical events. We view these two approaches as complementary: the stationary Markov framework is well-suited for studying unanticipated, recurrent events such as business cycles, whereas the EFP method is designed to analyze anticipated, nonrecurrent historical events, such as Brexit.

In special cases, nonstationary solutions produced by the EFP method can resemble stationary Markov solutions. For instance, a one-time, non-Markov parameter shift is analogous to a regime-switching Markov model with an absorbing state or to a news-shock Markov model with highly persistent (random walk) shocks. However, more complex EFP scenarios—comprising nonrecurrent periods of growth, transitions, shifts, or drifts—cannot be adequately captured by a single stationary Markov process. To illustrate, we compared the EFP solution with the news-shock method of Schmitt-Grohé and Uribe (2012) in the context of the bToTEM model and found that the differences between the two approaches were economically significant. While it may be possible to approximate complex non-Markov EFP scenarios by combining a sequence of stationary Markov solutions, it remains unclear whether such constructions would be computationally tractable.

The remainder of the paper is organized as follows: Section 2 provides an overview of the EFP methodology using a neoclassical optimal growth model as an example; Section 3 introduces a large-scale central banking bToTEM model; Section 4 presents five policy experiments; Section 5 compares the solutions under Markov and non-Markov news shocks; and finally, Section 6 concludes the paper.

1. ANTICIPATED NON-MARKOV EVENTS AND THE EFP METHOD AT A GLANCE

We use an example of a stylized neoclassical optimal growth model to illustrate the perturbation-based EFP methodology at a glance and to show how anticipated non-Markov events affect the macroeconomy.

1.1 *Anticipated Events*

Anticipated events refer to occurrences that an individual or group expects to take place in the future. In our analysis, economic agents may become aware of future changes affecting the economy's fundamentals (preferences, technology, economic environment, equations, variables, parameters, expectations) or policies (legislative, monetary, fiscal, social security). These future events can be deterministic or stochastic, with dates that may be fixed or uncertain. Importantly, the events we consider are non-Markov, meaning they are nonrecurrent and time-contingent (i.e., occurring at specific dates). Our approach contrasts with mainstream literature, which focuses

on Markov events that are recurrent and state-contingent. An example of an anticipated, non-Markov event is Brexit, which took place on January 31, 2020 following a referendum held in the UK on June 23, 2016. Once Brexit was announced, rational agents adjusted their behavior to optimally respond to the new environment before the event actually took place. In contrast, an example of recurrent, unanticipated Markov events commonly studied in the literature is the business cycle. Business cycles occur at random dates with state-contingent probabilities. Rational agents respond to business cycles by adapting their decision rules to the presence of uncertainty, such as by engaging in precautionary savings. Below, we illustrate the difference between unanticipated Markov and anticipated non-Markov events using the example of a stylized optimal growth model.

1.2 A Growth Model with Markov and Non-Markov Shocks

We consider a stylized stochastic growth model

$$\max E_0 \left[\sum_{t=0}^T \beta^t u(c_t) \right] \quad (1)$$

$$\text{s.t. } c_t + k_{t+1} = (1 - \delta)k_t + z_t f(k_t), \quad (2)$$

where c_t and k_t are consumption and capital, respectively; the initial condition (k_0, z_0) is given; the utility and production functions $u : \mathbb{R}_+ \rightarrow \mathbb{R}$ and $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ are twice continuously differentiable, strictly increasing, strictly quasi-concave, and satisfy the Inada conditions; z_t is the exogenous productivity level; $\beta \in (0, 1)$ is a discount factor; $\delta \in (0, 1]$ is a depreciation rate; $E_t[\cdot]$ is an expectation operator conditional on the t -period information set $\mathcal{F}_t$; and lifetime utility (1) is bounded.

We assume that the time horizon is infinite, $T = \infty$. A feasible path is a pair of random processes $\{c_t, k_{t+1}\}_{t=0}^{\infty}$ such that, given information sets $(\mathcal{F}_0, \mathcal{F}_1, \dots)$, they satisfy $c_t \geq 0$, $k_{t+1} \geq 0$, and (2) for $t \geq 0$. We assume that the set of feasible paths is not empty.

An optimal path is a feasible path that maximizes (1). We focus on the case where the optimal program of (1)-(2) is interior and satisfies the Euler equation:

$$u'(c_t) = \beta E_t[u'(c_{t+1})(1 - \delta + z_{t+1}f'(k_{t+1}))]. \quad (3)$$

1.3 A Nonstationary Model with Anticipated Events

For the sake of illustration, we assume that z_t in (ar1) follows the usual first-order autoregressive process with an additional term A_t that represents anticipated events:

$$\ln z_t = A_t + \rho \ln z_{t-1} + \sigma \varepsilon_t, \quad (4)$$

where $\sigma > 0$ and $|\rho| \leq 1$. Here, ε_t is a Markov shock defined by a probability distribution $\varepsilon_t \sim N(0, 1)$; A_t is a non-Markov process defined by a time-path $(A_0, A_1, \dots)$. The impact of anticipated events depends on what information about the sequence $(A_0, A_1, \dots)$ is provided to the agent, and when.

In the absence of anticipated events (i.e., $A_t = 0$ for all t), the model (1), (2), and (4) is a business cycle model whose solution is given by a time-invariant (stationary) Markov decision function $k_{t+1} = K^*(k_t, z_t)$ satisfying a fixed-point property. If we substitute the $t + 1$ -period function $K_{t+1}(k_{t+1}, z_{t+1})$ into the right side of (3) and compute the t -period function $K_t(k_t, z_t)$ on the left side of (3), we obtain the same fixed-point function $K_{t+1} = K_t = K^*$.¹

However, in the presence of non-Markov anticipated shocks, the decision rules are generally nonstationary (time-dependent). Intuitively, if at $t = 0$ the agent is informed of a shock $A_t = 1$ at $t \geq 1$, she reacts immediately by adjusting her decision rules $K_0^*, \dots, K_t^*$ to take advantage of a future productivity change. Her decision functions change over time; they depend on how many periods remain until the shock actually happens, meaning $K_0^* \neq \dots \neq K_t^*$.

We can consider a variety of non-Markov scenarios designed to represent the historical paths of actual economies. In particular, the economy may face multiple anticipated events—say, in periods $t = 10$, $t = 17$, and $t = 23$. These events may happen with specific probabilities, such as 50%, 30%, and 10%. Anticipated shocks may affect multiple equations and parameters and may be correlated with, or conditional on, one another. Furthermore, this information can be revealed to the agent all at once or gradually.

For stationary models, we only need to construct one fixed-point decision rule $K_0^* = K_1^* = \dots = K^*$ that corresponds to a given Markov process, as this rule works for all time periods. Conventional global nonlinear solution methods iterate backward on Euler equations until a fixed-point decision rule K^* is found. Similarly, conventional perturbation methods construct Taylor expansions by restricting the decision rules to be identical at t and $t + 1$. However, these methods fail for the nonstationary models we study because fixed-point decision rules either do not exist or are suboptimal in the presence of non-Markov shocks. For our nonstationary model, we must construct an infinite sequence of optimal decision functions $K_0^*, K_1^*, \dots$.

1.4 Turnpike Theorem

In the present paper, we develop a perturbation solution framework for approximating a sequence of decision functions $K_0^*, K_1^*, \dots$ in the presence of anticipated non-Markov shocks for a broad and empirically relevant class of nonstationary models that satisfy the turnpike theorem. The turnpike theorem states that the trajectory of a finite-horizon economy converges to that of an infinite-horizon economy as the time horizon increases. Let us illustrate this theorem for the model (1), (2), and (4) under a parameterization that yields a closed-form solution. Specifically, $u(c) = \ln(c)$,

1. A Markov switching model (e.g., Davig and Leeper, 2007) and a recurrent news shock model (e.g., Schmitt-Grohé and Uribe, 2012) belong to the class of stationary models.




Fig. 1. Capital trajectories.

$f(k) = k^\alpha$, $A_t = \gamma_A^t$, and we set $\delta = 1$, $\beta = 0.99$, $\alpha = 0.36$, $\rho = 0.95$, $\sigma = 0.01$, and $\gamma_A = 1.01$.

Here, A_t is an anticipated deterministic technological progress (parameter drift), where $\gamma_A \geq 1$ is a growth rate. Let us compare the trajectories of the finite- and infinite-horizon economies. In a finite-horizon model where life ends at $T + 1$, we have $k_{T+1} = 0$, and the decision rules for the previous periods $T, T - 1, \dots$ are given by:

$$\begin{aligned} k_T &= \left(\frac{\alpha\beta}{1 + \alpha\beta} \right) z_{T-1} k_{T-1}^\alpha A_{T-1}, \quad k_{T-1} \\ &= \left(\frac{\alpha\beta(1 + \alpha\beta)}{1 + \alpha\beta(1 + \alpha\beta)} \right) z_{T-2} k_{T-2}^\alpha A_{T-2}, \quad \text{etc.} \end{aligned} \quad (5)$$

By extending the time horizon T to infinity, we obtain the solution for the infinite-horizon economy. In Figure 1, we plot the capital series trajectories of economies with finite horizons of $T = 15$ and $T = 25$ under identical realizations of shocks, alongside the trajectory of the infinite-horizon economy ($T = \infty$). From the figure, we observe the following: if all three economies start with the same initial capital stock, they follow a virtually identical path for some time and diverge only in close proximity to the terminal date. Thus, during the initial periods τ , the infinite-horizon solution can be accurately approximated by the finite-horizon solution. This observation illustrates the essence of the turnpike theorem; see MMTT (2020) for a formal discussion. We now have a foundation for our nonstationary analysis. Since a finite-

horizon model is solved by backward iteration from a given terminal condition K_T^* , producing a sequence of time-varying optimal decision rules $K_{T-1}^*, K_{T-2}^*, \dots, K_0^*$, we can introduce any anticipated non-Markov process that combines parameter shocks, shifts, and drifts.

1.5 EFP Method

On the basis of the turnpike theorem, MMTT (2020) introduce a numerical method—the EFP—for approximating a sequence of optimal decision functions in an infinite-horizon economy $(K_0^*, K_1^*, \dots)$ by a finite-horizon solution $(K_0, \dots, K_\tau)$ during the initial τ periods, see Algorithm 1 below:

Algorithm 1: Extended function path (EFP)

Step 1: Terminal condition. Choose $T \gg \tau$ and construct terminal condition K_T .

Step 2: Backward iteration. Given the terminal condition K_T , iterate backward on the Euler equation to construct a path of $(K_{T-1}, \dots, K_0)$.

Step 3: Turnpike property. Verify that the initial τ functions $(K_0, \dots, K_\tau)$ are not sensitive to the choice of time horizon T and terminal condition K_T .

Use $(K_0, \dots, K_\tau)$ as an approximate solution and discard the remaining $(K_{\tau+1}, \dots, K_T)$.

MMTT (2020) use the EFP method to analyze several challenging applications with time-dependent decision rules, including unbalanced stochastic growth models, entry into and exit from a monetary union, information news, anticipated policy regime switches, and deterministic seasons. A shortcoming of the proposed EFP method is its reliance on global projection-style techniques, which entail a high computational expense. To mitigate this, MMTT (2020) employ Smolyak (sparse) grids, but the method remains costly; it takes several minutes to implement even for the simplest optimal growth model.

1.6 Perturbation EFP Method

In this paper, we analyze time-dependent scenarios in large-scale New Keynesian models of central banking. The global EFP method proposed by MMTT (2020) is prohibitively expensive, or even intractable, for such large models. Moreover, implementing global solutions requires significant coding effort, which becomes particularly challenging for scale. To address these limitations, we introduce an alternative version of the EFP method that leverages relatively inexpensive and widely used perturbation techniques. To ensure our analysis is accessible to a broader audience, we implement our approach using Dynare—a popular automated software for constructing perturbation solutions in economics (see <https://www.dynare.org>).

The built-in Dynare perturbation software is designed to construct a stationary solution to a Markov model. It uses the same deterministic steady state to expand the decision rules for periods t and $t + 1$, and it imposes a fixed-point restriction on the coefficients of the decision rules to remain identical in periods t and $t + 1$. However,




Fig. 2. Anticipated events.

both of these restrictions are invalid for our nonstationary analysis, making Dynare's built-in solver unsuitable for our purposes.

We design an EFP perturbation framework that builds upon Dynare routines to address three key points: (i) we solve for a time-dependent sequence (path) of states around which the perturbation solutions are constructed—these points replace the deterministic steady state; (ii) we allow the decision rules to differ in consecutive periods—this construction replaces the standard fixed-point solution; and (iii) we modify the solver to construct the optimal sequence of decision rules instead of finding a fixed-point limiting solution. In the remainder of the section, we illustrate the solutions constructed by the EFP perturbation method.

In some nonstationary models, identifying the point(s) around which to construct the Taylor expansion is straightforward. For instance, in our policy experiments with the bToTEM model, a shift in parameters causes a change in the steady state. In these cases, we construct perturbation solutions around the old and new steady states, assuming the economy switches between the solutions when the parameters change. However, this approach is not feasible in models with unbalanced growth. For such models, we first construct a deterministic growth path that the economy would follow in the absence of business-cycle shocks. This growth path serves as a set of points around which the Taylor expansion is constructed. For a detailed example, see online Appendix A.²

Illustration 1: Anticipated events with uncertain outcomes. In the first experiment shown in Figure 2, an unanticipated productivity increase at $t = 1$ (solid blue line) is compared with an anticipated productivity increase that is announced at $t = 1$ but occurs either 1 year later (at $t = 5$, dashed red line) or 2 years later (at $t = 9$, dash-

2. Some nonstationary models admit variable scaling to render a scaled model with a deterministic steady state. A well-known example is a growth model with labor-augmenting technological progress alongside homothetic preferences and technology, which follows a balanced growth path and can thus be converted into a stationary model; see King, Plosser, and Rebelo (1988). However, our focus is not on these exceptional cases, but rather on more general, nonstationary scenarios.




Fig. 3. Anticipated uncertain outcome.

dotted green line). We observe well-pronounced anticipatory effects on the red and green lines (the variables shown are in percentage deviations from the risky steady state). Immediately after the productivity growth is announced, the agent takes advantage of the future productivity increase by reducing (or “eating up”) her capital to increase consumption. The instantaneous increase in consumption is larger for the earlier productivity increase because, in this case, the additional resources are spread over a shorter horizon.

Illustration 2: Anticipated events with uncertain outcomes. In the second experiment shown in Figure 3, we consider the case where the agent is informed at $t = 1$ that a productivity increase will occur in 1 year (period $t = 5$). We compare the scenario where the productivity increase is certain (dash-dotted green line) with the scenario where it occurs with a 50 probability (the solid blue and dashed red lines show the paths when the increase actually occurs and does not occur, respectively). We observe that the anticipated growth in consumption is much larger for a productivity increase that occurs with certainty (green line) than for uncertain ones (red and blue lines). When the productivity increase is realized and the uncertainty is resolved, we see an additional boost in consumption for the economy where productivity grew, and a sharp decline in consumption for the economy where it did not.

Illustration 3: Modeling a historical path with anticipated events. In the final experiment shown in Figure 4, we consider an economy facing a sequence of anticipated events (solid blue line), including a temporary positive shock, a prolonged but temporary negative regime shift, and a permanent positive regime shift. Specifically, our goal is to illustrate how our framework allows us to model a mixture of certain and uncertain events to reflect the historical trajectory of an actual economy. For comparison, we also show a naive solution (dashed red line) wherein future events are not anticipated; at each moment, the agent believes that the current state of the economy will persist forever and perceives each productivity change as a surprise. This solution is computed in a point-by-point manner by constructing the usual perturbation around the current state. We observe that the anticipatory effects can be quite sizable.




Fig. 4. Example path.

The EFP framework can also be used to model how economic agents learn about the importance of future events. Learning has not been addressed in our analysis thus far. Our primary goal is to study short-term boosts in economic activity following the announcements of economic changes. However, certain announcements can serve another purpose: to ensure that households, firms, and markets have sufficient time to understand and prepare for future changes. This is particularly true for significant shifts in policy frameworks, such as the adoption of PLT or AIT, as discussed in Section 4.

There are different ways to model how agents learn about the importance of future events. One possibility is to assume that the agent initially underestimates the shock but gradually perceives its true size. For example, in period $t = 0$, the authority announces that at $t = 10$, there will be a shock of unit size $A_{10} = 1$. However, the agent believes the size of the shock will be smaller; specifically, in periods $t = 0, 1, 2, \dots, 10$, the agent expects A_{10} to take values $0, 1/10, 2/10, \dots, 1$, respectively, recognizing the correct size of the shock only at the end. Another possibility is to model learning in terms of the probability of a shock, assuming the agent initially underestimates the probability but improves her forecast over time. For example, in periods $t = 0, 1, 2, \dots, 10$, the agent understands that $A_{10} = 1$ will happen with probabilities $0, 1/10, 2/10, \dots, 1$, recognizing the correct probability only at the end. The EFP framework allows us to model these and other learning schemes in order to accurately represent the historical processes taking place in actual economies.

2. THE LARGE-SCALE CENTRAL BANKING MODEL

Today, central banks, leading international organizations, and government agencies use large-scale macroeconomic models for projection and policy analysis. A prominent example is the ToTEM employed by the Bank of Canada. This model includes heterogeneous utility-maximizing consumers, multiple profit-maximizing production

sectors, monetary and fiscal authorities, and a foreign sector. ToTEM is highly complex; see Corrigan et al. (2021) and Dorich et al. (2013).

In this paper, we consider a scaled-down version of ToTEM developed by LMM (2020). Like the full-scale model, the “baby” ToTEM (hereafter, bToTEM) is a small open-economy model. It features multiple New Keynesian Phillips curves—one driven by sticky prices in domestic production, one by sticky wages, and one by sticky import prices. We incorporate rule-of-thumb price setters in line with Galí and Gertler (1999). We assume quadratic adjustment costs for investment and convex costs for capital utilization to generate more realistic model dynamics, particularly concerning monetary policy transmission. International trade consists of exporting domestic consumption goods and commodities, alongside importing foreign goods for domestic production. Even though bToTEM is significantly smaller than ToTEM (comprising only 47 equations and unknowns, including 21 state variables), it generates realistic impulse response functions for the Canadian economy that closely mirror those produced by the full-scale ToTEM model; see LMM (2020) for detailed comparisons. In this section, we present the bToTEM central banking model and the implementation of the EFP analysis.

Final-good production. Final consumption goods are produced in two stages. In the first stage, intermediate goods are produced competitively using labor, capital, commodities, and imports. In the second stage, final goods are aggregated from differentiated goods, each produced by a monopolistically competitive firm using intermediate goods and other final goods. These final goods can either be consumed by households or transformed via linear technologies into other types of goods, specifically investment goods and noncommodity exports.

In the first production stage, a representative, perfectly competitive firm produces an intermediate good by solving the following profit maximization problem:

$$\begin{aligned} & \max_{\{Z_t^g, Z_t^n, L_t, K_t, I_t, COM_t^d, M_t, u_t, d_t\}} E_0 \sum_{t=0}^{\infty} \mathcal{R}_{0,t} (P_t^c Z_t^n - W_t L_t - P_t^i I_t - P_t^{com} COM_t^d - P_t^m M_t) \\ \text{s.t. } Z_t^g &= \left[\delta_l (A_t L_t)^{\frac{\sigma-1}{\sigma}} + \delta_k (u_t K_{t-1})^{\frac{\sigma-1}{\sigma}} + \delta_{com} (COM_t^d)^{\frac{\sigma-1}{\sigma}} + \delta_m (M_t)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}, \end{aligned} \quad (6)$$

$$\log A_t = \varphi_a \log A_{t-1} + (1 - \varphi_a) \log \bar{A} + \xi_t^a, \quad (7)$$

$$K_t = (1 - d_t) K_{t-1} + I_t, \quad (8)$$

$$d_t = d_0 + \bar{d} e^{\rho(u_t - 1)}, \quad (9)$$

$$Z_t^n = Z_t^g - \frac{\chi_i}{2} \left(\frac{I_t}{I_{t-1}} - 1 \right)^2 I_t, \quad (10)$$

where Z_t^g and Z_t^n are gross production and net (of adjustment costs) production of final goods; L_t , K_t , I_t , COM_t^d , M_t , u_t , and d_t are labor, capital, investment, commodity inputs, imports, capital utilization, and depreciation rate, respectively; A_t is the level of labor-augmenting technology; ξ_t^a is a normally distributed variable, and φ_a is an autocorrelation coefficient. The firm discounts nominal payoffs according to household's stochastic discount factor $\mathcal{R}_{t,t+j} = \beta^j (\lambda_{t+j}/\lambda_t) (P_t/P_{t+j})$, where λ_t is the household's marginal utility of consumption, and P_t is the final good price. Investment goods and noncommodity exports are assumed to be produced from the final goods according to linear technology, $P_t^i = \iota_i P_t$ and $P_t^{nc} = \iota_x P_t$, where P_t^i and P_t^{nc} are the price of investment goods and noncommodity exports goods, respectively.

In the second stage of production, monopolistically competitive firms produce a continuum of differentiated goods. These differentiated goods are then aggregated into a single final good by an aggregating firm that solves the following cost minimization problem:

$$\begin{aligned} \min_{\{Z_{it}\}} \int_0^1 P_{it} Z_{it} di \\ \text{s.t. } Z_t = \left(\int_0^1 Z_{it}^{\frac{\varepsilon-1}{\varepsilon}} di \right)^{\frac{\varepsilon}{\varepsilon-1}}, \end{aligned}$$

where Z_t and P_{it} are given; Z_{it} is a differentiated good i . The cost minimization implies the following demand function for the differentiated good i :

$$Z_{it} = \left(\frac{P_{it}}{P_t} \right)^{-\varepsilon} Z_t, \quad \text{with } P_t \equiv \left(\int_0^1 P_{it}^{1-\varepsilon} di \right)^{\frac{1}{1-\varepsilon}}.$$

Each differentiated good is produced from the intermediate goods and from the final goods using technology featuring perfect complementarity,

$$Z_{it} = \min \left(\frac{Z_{it}^n}{1 - s_m}, \frac{Z_{it}^{mi}}{s_m} \right), \quad (11)$$

where Z_{it}^n is an intermediate good and Z_{it}^{mi} is a final good input, and s_m is a Leontief parameter.

There are two types of the monopolistically competitive firms producing differentiated goods: rule-of-thumb firms of measure ω and forward-looking firms of measure $1 - \omega$. Both rule-of-thumb firms and forward-looking firms index their price to the inflation target $\bar{\pi}_t$ with probability θ as $P_{it} = \bar{\pi}_t P_{i,t-1}$. The rule-of-thumb firms par-

tially index their price to lagged inflation and target inflation with probability $1 - \theta$,

$$P_{it} = (\pi_{t-1})^\gamma (\bar{\pi}_t)^{1-\gamma} P_{i,t-1}. \quad (12)$$

Forward-looking firms choose their prices P_t^* with probability $1 - \theta$ to maximize profits generated when the price remains effective

$$\begin{aligned} \max_{P_t^*} E_t \sum_{j=0}^{\infty} \theta^j \mathcal{R}_{t,t+j} & \left(\prod_{k=1}^j \bar{\pi}_{t+k} P_t^* Z_{i,t+j} - (1 - s_m) P_{t+j}^z Z_{i,t+j} - s_m P_{t+j} Z_{i,t+j} \right) \\ \text{s.t. } Z_{i,t+j} & = \left(\frac{\prod_{k=1}^j \bar{\pi}_{t+k} P_t^*}{P_{t+j}} \right)^{-\varepsilon} Z_{t+j}. \end{aligned} \quad (13)$$

The production in the first stage Z_t^n and that in the second stages Z_t are related via price dispersion Δ_t ,

$$Z_t^n = \int_0^1 Z_{it}^n di = (1 - s_m) \int_0^1 Z_{it} di = (1 - s_m) \int_0^1 \left(\frac{P_{it}}{P_t} \right)^{-\varepsilon} Z_t di = (1 - s_m) \Delta_t Z_t, \quad (14)$$

where $\Delta_t \equiv \int_0^1 \left(\frac{P_{it}}{P_t} \right)^{-\varepsilon} di$.

Production of commodities. Commodities are produced by a domestic firm using final goods and land as inputs. They are sold domestically or exported to the rest of the world. The domestic firm solves

$$\begin{aligned} \max_{Z_t^{com}, COM_t} & \{ P_t^{com} COM_t - P_t Z_t^{com} \} \\ \text{s.t. } COM_t & = (Z_t^{com})^{s_z} (A_t F)^{1-s_z} - \frac{\chi_{com}}{2} \left(\frac{Z_t^{com}}{Z_{t-1}^{com}} - 1 \right)^2 Z_t^{com}, \end{aligned} \quad (15)$$

where Z_t^{com} is the final good input, and F is a fixed production factor, which may be considered as land. Similar to production of final goods, the commodity producers incur quadratic adjustment costs when they adjust the level of final good input. The commodities are sold domestically or exported to the rest of the world, $COM_t = COM_t^d + X_t^{com}$. They are sold at the world price adjusted by the nominal exchange rate, $P_t^{com} = e_t P_t^{comf}$, where e_t is the nominal exchange rate (i.e., domestic price of a unit of foreign currency), and P_t^{comf} is the world commodity price; in real terms, the latter price is given by $p_t^{com} = s_t p_t^{comf}$, where $p_t^{com} \equiv P_t^{com} / P_t$ and $p_t^{comf} \equiv P_t^{comf} / P_t^f$ are domestic and foreign relative prices of commodities, respectively, P_t^f is the foreign consumption price level, and $s_t = e_t P_t^f / P_t$ is the real exchange rate.

Production of imports. The representative perfectly competitive firm produces the final imported good M_t from a continuum of intermediate imported goods M_{it} and solves the following cost-minimization problem,

$$\begin{aligned} \min_{\{M_{it}\}} & \int_0^1 P_{it}^m M_{it} di \\ \text{s.t. } M_t &= \left(\int_0^1 M_{it}^{\frac{\varepsilon_m-1}{\varepsilon_m}} di \right)^{\frac{\varepsilon_m}{\varepsilon_m-1}}, \end{aligned}$$

where M_{it} is an intermediate imported good i . The demand for an intermediate imported good i is given by

$$M_{it} = \left(\frac{P_{it}^m}{P_t^m} \right)^{-\varepsilon_m} M_t, \quad \text{with } P_t^m \equiv \left(\int_0^1 (P_{it}^m)^{1-\varepsilon_m} di \right)^{\frac{1}{1-\varepsilon_m}}.$$

Prices of the intermediate imported goods are sticky in a similar way as the prices of the differentiated final goods. A measure ω_m of the importers follows the rule-of-thumb pricing, and the others are forward looking. The optimizing forward-looking importers choose the price P_t^{m*} in order to maximize profits generated when the price remains effective

$$\begin{aligned} \max_{P_t^{m*}} E_t \sum_{j=0}^{\infty} (\theta_m)^j \mathcal{R}_{t,t+j} & \left(\prod_{k=1}^j \bar{\pi}_{t+k} P_t^{m*} M_{i,t+j} - e_{t+j} P_{t+j}^{mf} M_{i,t+j} \right) \\ M_{i,t+j} &= \left(\frac{\prod_{k=1}^j \bar{\pi}_{t+k} P_t^{m*}}{P_{t+j}^m} \right)^{-\varepsilon_m} M_{t+j}, \end{aligned}$$

where P_t^{mf} is the price of imports in the foreign currency. All importers face the same marginal cost given by the foreign price of imports.

Households. Households maximize the lifetime utility by choosing holdings of domestic and foreign-currency denominated bonds, labor, and consumption, and they are subject to habits in consumption. Each household supplies a variety of differentiated labor service to the labor market, which is monopolistically competitive. The representative household of type h solves the following utility-maximization problem:

$$\begin{aligned} \max_{C_t, L_{ht}, B_t, B_t^f} E_0 \sum_{t=0}^{\infty} \beta^t & \left\{ \frac{\mu}{\mu-1} (C_t - \xi \bar{C}_{t-1})^{\frac{\mu-1}{\mu}} \exp \left(\frac{\eta(1-\mu)}{\mu(1+\eta)} \int_0^1 (L_{ht})^{\frac{\eta+1}{\eta}} dh \right) \eta_t^c \right\}, \\ \text{s.t. } P_t C_t + \frac{B_t}{R_t} + \frac{e_t B_t^f}{R_t^f (1 + \kappa_t^f)} &= B_{t-1} + e_t B_{t-1}^f + \int_0^1 W_{ht} L_{ht} dh + \Pi_t, \end{aligned} \quad (16)$$

$$\log \eta_t^c = \varphi_c \log \eta_{t-1}^c + \xi_t^c, \quad (17)$$

where C_t, L_{ht}, B_t, B_t^f are consumption of final goods, labor service of type h , holdings of domestic and foreign-currency denominated bonds, respectively; $\bar{C}_t$ is the aggregate consumption, taken by the household as given; $\beta \in (0, 1)$ is a subjective discount factor; μ and η are the utility-function parameters; η_t^c is a consumption demand shock, ξ_t^c is a normally distributed variable, and φ_c is an autocorrelation coefficient; R_t and R_t^f are domestic and foreign nominal interest rate, respectively; κ_t^f is the risk premium on the foreign interest rate; W_{ht} is the nominal wage of labor of type h ; Π_t are profits paid by the firms. The representative household supplies a variety of differentiated labor service to the labor market, which is monopolistically competitive.

Labor packer. A labor packer aggregates differentiated labor services by solving

$$\begin{aligned} \min_{\{L_{ht}\}} & \int_0^1 W_{ht} L_{ht} dh \\ \text{s.t. } L_t &= \left(\int_0^1 L_{ht}^{\frac{\varepsilon_w - 1}{\varepsilon_w}} dh \right)^{\frac{\varepsilon_w}{\varepsilon_w - 1}}, \end{aligned}$$

where L_t is aggregated labor demanded by firms in the first stage of production. Cost minimization of the labor packer implies the following demand for individual labor:

$$L_{ht} = \left(\frac{W_{ht}}{W_t} \right)^{-\varepsilon_w} L_t, \quad \text{with } W_t \equiv \left(\int_0^1 W_{ht}^{1-\varepsilon_w} dh \right)^{\frac{1}{1-\varepsilon_w}}. \quad (18)$$

Labor unions. Labor unions set wages. There are two types of labor unions: rule-of-thumb unions of measure ω_w and forward-looking unions of measure $1 - \omega_w$. Within each type, with probability θ_w , the labor unions index their wage to the inflation target $\bar{\pi}_t$ as follows $W_{it} = \bar{\pi}_t W_{i,t-1}$. The rule-of-thumb unions that do not index their wage in the current period follow the rule

$$W_{it} = (\pi_{t-1}^w)^{\gamma_w} (\bar{\pi}_t)^{1-\gamma_w} W_{i,t-1}. \quad (19)$$

A forward-looking unions that do not index its wage solves

$$\max_{W_t^*} E_t \sum_{j=0}^{\infty} (\beta \theta_w)^j \left\{ \frac{\mu}{\mu - 1} (C_{t+j} - \xi \bar{C}_{t+j-1})^{\frac{\mu-1}{\mu}} \exp \left(\frac{\eta(1-\mu)}{\mu(1+\eta)} \int_0^1 (L_{ht+j})^{\frac{\eta+1}{\eta}} dh \right) \eta_{t+j}^c \right\} \quad (20)$$

$$\text{s.t. } L_{h,t+j} = \left(\frac{\prod_{k=1}^j \bar{\pi}_{t+k} W_t^*}{W_{t+j}} \right)^{-\varepsilon_w} L_{t+j}, \quad (21)$$

$$P_{t+j}C_{t+j} = \prod_{k=1}^j \bar{\pi}_{t+k} W_t^* L_{h,t+j} dh + \Psi_{t+j}, \quad (22)$$

where Ψ_{t+j} includes terms in budget constraints (16) other than C_{t+j} and $L_{h,t+j}$.

Monetary authority. The central bank uses a Taylor rule to set the short-term nominal interest rate,

$$R_t = \rho_r R_{t-1} + (1 - \rho_r) [\bar{R}_t + \rho_\pi (\pi_t - \bar{\pi}_t) + \rho_Y (\log Y_t - \log \bar{Y}_t)] + \eta_t^r, \quad (23)$$

where ρ_r measures the degree of smoothing of the interest rate; $\bar{R}_t$ is the nominal neutral interest rate; ρ_π measures a response to the inflation gap; $\bar{\pi}_t$ is the inflation target; ρ_Y measures a response to the output gap; $\bar{Y}_t$ is potential output; η_t^r is an interest rate shock following a process

$$\eta_t^r = \varphi_r \eta_{t-1}^r + \xi_t^r,$$

where ξ_t^r is a normally distributed variable, and φ_r is an autocorrelation coefficient. Potential output changes with productivity according to

$$\log \bar{Y}_t = \varphi_z \log \bar{Y}_{t-1} + (1 - \varphi_z) \log \left(\frac{A_t \bar{Y}}{\bar{A}} \right).$$

If an effective lower bound (ELB) R_t^{elb} is imposed on the nominal interest rate, the interest rate is determined as a maximum of (23) and R_t^{elb} :

$$R_t = \max \{ R_t^{elb}, \Phi_t \}.$$

The ELB restriction was never binding in our experiments. However, our earlier LMM (2020) paper studied a stationary version of bToTEM model with an occasionally binding ELB constraint using IRIS perturbation software (see <https://iris.igpmn.org>), and Occbin perturbation toolbox of Guerrieri and Iacoviello (2015); these methods can also be used in the context of our nonstationary analysis.

Foreign demand for noncommodity exports. By analogy with the demand for imports, the foreign demand function for noncommodity exports is assumed to be

$$X_t^{nc} = \gamma^f \left(\frac{P_t^{nc}}{e_t P_t^f} \right)^{-\phi} Z_t^f, \quad (24)$$

where P_t^{nc} is a domestic price of noncommodity exports; γ^f is the demand-function parameter. In real terms, we have

$$X_t^{nc} = \gamma^f \left(\frac{s_t}{p_t^{nc}} \right)^\phi Z_t^f. \quad (25)$$

Balance of payments. The balance of payments in nominal terms is given by

$$\frac{e_t B_t^f}{R_t^f (1 + \kappa_t^f)} - e_t B_{t-1}^f = P_t^{nc} X_t^{nc} + P_t^{com} X_t^{com} - P_t^m M_t, \quad (26)$$

where B_t^f is domestic holdings of foreign-currency denominated bonds, and R_t^f is the nominal interest rate on the bonds. By normalizing the bonds holdings as $b_t^f \equiv \frac{e_t B_t^f}{\pi_{t+1}^f P_t \bar{Y}}$, the balance of payments in real terms becomes

$$\frac{b_t^f}{r_t^f (1 + \kappa_t^f)} - b_{t-1}^f \frac{s_t}{s_{t-1}} = \frac{1}{\bar{Y}} (p_t^{nc} X_t^{nc} + p_t^{com} X_t^{com} - p_t^m M_t), \quad (27)$$

where r_t^f is the real interest rate on the foreign-currency denominated bonds.

Rest-of-the-world economy. The rest of the world is specified by three exogenous processes that, respectively, describe the evolution of foreign output Z_t^f , the foreign real interest rate r_t^f , and the foreign commodity price p_t^{comf} ,

$$\log Z_t^f = \varphi_{zf} \log Z_{t-1}^f + (1 - \varphi_{zf}) \log \bar{Z}^f + \xi_t^{zf}, \quad (28)$$

$$\log r_t^f = \varphi_{rf} \log r_{t-1}^f + (1 - \varphi_{rf}) \log \bar{r} + \xi_t^{rf}, \quad (29)$$

$$\log p_t^{comf} = \varphi_{comf} \log p_{t-1}^{comf} + (1 - \varphi_{comf}) \log \bar{p}^{comf} + \xi_t^{comf}, \quad (30)$$

where ξ_t^{zf} , ξ_t^{rf} , and ξ_t^{comf} are normally distributed random variables, and φ_{zf} , φ_{rf} , and φ_{comf} are autocorrelation coefficients.

Uncovered interest rate parity. We impose an augmented uncovered interest rate parity condition

$$e_t = E_t \left[(e_{t-1})^\varkappa \left(e_{t+1} \frac{R_t^f (1 + \kappa_t^f)}{R_t} \right)^{1-\varkappa} \right]. \quad (31)$$

Stationarity condition for the open-economy model. The risk premium κ_t^f is a decreasing function of foreign assets

$$\kappa_t^f = \varsigma (\bar{b}^f - b_t^f), \quad (32)$$

where $\bar{b}^f$ is the steady state level of the normalized bond holdings. This assumption ensures a decreasing rate of return to foreign assets.

Summary of the model's variables. For each period t , there are the following four types of variables in this model: 47 endogenous (or non-predetermined) variables,

$$y_t \equiv \left\{ \begin{array}{c} F_{1t}, F_{2t}, F_{1t}^w, F_{2t}^w, F_{1t}^m, F_{2t}^m, q_t, \lambda_t, s_t, \\ L_t, K_t, I_t, COM_t^d, M_t, u_t, d_t, Z_t^g, Z_t^z, Z_t, C_t, Y_t, \pi_t, rmc_t, \Delta_t, \pi_t^m, \bar{\pi}_t, p_t^m, R_t, p_t^z, w_t, \\ MPK_t, R_t^k, p_t^j, \kappa_t^f, b_t^f, X_t^{nc}, X_t^{com}, COM_t, Z_t^{com}, \pi_t^w, w_t^*, \Delta_t^w, \bar{Y}_t, p_t^{com}, p_t^{nc}, p_t^{mf}, p_t^y, \end{array} \right\},$$

where $\{F_{1t}, F_{2t}\}$, $\{F_{1t}^w, F_{2t}^w\}$, $\{F_{1t}^m, F_{2t}^m\}$ are supplementary variables in Phillips curves for prices, wages, and imports, respectively; q_t is Tobin's q ; rmc_t and MPK_t are real marginal cost and marginal productivity of capital, respectively; p_t^{nc} and p_t^y are prices of noncommodity goods and output, respectively. We have 15 endogenous state variables:

$$\{C_{t-1}, R_{t-1}, s_{t-1}, \pi_{t-1}, \Delta_{t-1}, w_{t-1}, \pi_{t-1}^w, \Delta_{t-1}^w, p_{t-1}^m, \pi_{t-1}^m, I_{t-1}, Z_{t-1}^{com}, b_{t-1}^f, \bar{Y}_{t-1}, K_{t-1}\},$$

where π_{t-1}^w , Δ_{t-1}^w , p_{t-1}^m , π_{t-1}^m are wage inflation, wage dispersion, and price and inflation of imports; we have 19 endogenous forward variables:

$$\left\{ \begin{array}{c} F_{1t+1}, F_{2t+1}, F_{1t+1}^w, F_{2t+1}^w, F_{1t+1}^m, F_{2t+1}^m, \lambda_{t+1}, q_{t+1}, u_{t+1}, I_{t+1}, \\ Z_{t+1}^{com}, \pi_{t+1}, \bar{\pi}_{t+1}, \pi_{t+1}^m, \pi_{t+1}^w, p_{t+1}^{com}, p_{t+1}^z, MPK_{t+1}, s_{t+1} \end{array} \right\},$$

and we have six exogenous state variables:

$$z_t \equiv \{A_t, \eta_t^R, \eta_t^c, p_t^{comf}, r_t^f, Z_t^f\}.$$

Our calibration procedure for bToTEM is described in online Appendix C, and it closely follows the calibration of the full-scale ToTEM model.

The implementation of the perturbation EFP method for bToTEM model. LMM (2020) construct conventional time-invariant solutions to the stationary bToTEM model.³ That paper finds that high-order perturbation solutions are sufficiently accurate in the bToTEM model. Since our nonstationary analysis is more computationally costly, we limit our attention to perturbation solutions only. Here, we construct novel time-dependent (nonstationary) solutions to the model. We model announcements about economic policies that will be implemented at future dates as anticipated non-Markov news shocks, and we analyze the reaction of economic agents to such announcements. Assuming that we are at $t = 0$ and that a given policy will be implemented at $T > 0$, we construct a sequence of optimal decision functions for periods $t = 0, 1, \dots, T$ that characterize the anticipatory effects. (Naturally, the optimal decision functions will depend on how far the economy is from the moment the policy

3. LMM (2020) compare perturbation solutions with more accurate global projection solutions constructed using deep learning analysis.

is introduced.) Below, we outline our perturbation-based framework for analyzing economies with non-Markov news shocks.

We represent the bToTEM model as an infinite-horizon nonstationary equilibrium problem in which a solution is characterized by a set of equilibrium conditions for $t = 0, 1, \dots$,

$$E_t[G_t(y_{t-1}, y_t, y_{t+1}, z_t, z_{t+1})] = 0, \quad (33)$$

$$z_{t+1} = Z_t(z_t, \epsilon_{t+1}), \quad (34)$$

where the initial condition (z_0, y_{-1}) is given; E_t denotes the expectations operator conditional on information available at t ; $z_t \in \mathbb{R}^{d_z}$ is a vector of exogenous state variables at t ; Z_t is a time-dependent law of motion for z_t ; $y_t \in \mathbb{R}^{d_y}$ is a vector of endogenous variables; $\epsilon_{t+1} \in \mathbb{R}^{d_\epsilon}$ is a vector of shocks; and G_t is a continuously differentiable vector function. Note that the latter function is time-dependent because the model is nonstationary (due to, e.g., time-dependent parameters in policy rules, the production function, or the utility function). A solution is given by a set of time-dependent equilibrium functions $y_t = Y_t(z_t, y_{t-1})$ that satisfy (33), (34) in the relevant area of the state space.

Our perturbation analysis proceeds in the following two steps:

Step I: Solving a T -period stationary economy. Assume that in a very remote period T , the economy becomes stationary, that is, $G_t(\cdot) = G(\cdot)$ and $Z_t(\cdot) = Z(\cdot)$ for all $t \geq T$. Therefore, the system (33) and (34) becomes:

$$E_t[G(y_{t-1}, y_t, y_{t+1}, z_t, z_{t+1})] = 0, \quad (35)$$

$$z_{t+1} = Z(z_t, \epsilon_{t+1}). \quad (36)$$

Solving (35), (36) allows us to find the terminal solution $y_T = Y_T(z_T, y_{T-1})$.

Step II: Constructing a function path. Using the T -period solution $y_T = Y_T(z_T, y_{T-1})$ as a terminal condition, we iterate backward for $T - 1, \dots, 1$ on the corresponding equilibrium conditions to construct a sequence (path) of time-dependent value and decision functions $\{Y_{T-1}(\cdot), \dots, Y_1(\cdot)\}$. For example, for period t , the system on which we iterate backward is:

$$E_t[G_t(y_{t-1}, y_t, Y_{t+1}(z_{t+1}, y_t), z_t, z_{t+1})] = 0,$$

$$z_{t+1} = Z_t(z_t, \epsilon_{t+1}).$$

Here, we solve for today's endogenous variables y_t , given tomorrow's functions $z_{t+1} = Z_t(z_t, \epsilon_{t+1})$ and $y_{t+1} = Y_{t+1}(z_{t+1}, y_t)$.

In both steps, we use perturbation to find numerical approximations of the decision functions. In Step I, a Taylor expansion of the policy functions in a stationary model is found around the deterministic steady state y of the model. In Step II, we consider

two alternative options. The first option is to find solutions for y_{t+1} and y_t around y_t and y_{t-1} , respectively, such that $y_t = y_{t-1} \equiv y$; in Dynare, this can be implemented by coding y_t and y_{t+1} using the same variable names. The other option is to consider y_{t+1} and y_t perturbed around y_t and y_{t-1} , respectively, such that $y_t = y_{t-1} \gamma_{y,t-1}$, where $\gamma_{y,t-1}$ is a time-dependent growth rate; in Dynare, this can be implemented by coding y_t and y_{t+1} with different variable names. For the bToTEM model, we limit our attention to the second option, but in online Appendix A, we compare both options using the example of a neoclassical optimal growth model.

3. ANALYZING ANTICIPATED NEWS SHOCKS

In this section, we present a series of policy experiments in which we consider anticipated changes in one or more of the model's parameters. In all figures, the variables are shown as percentage deviations from the initial risky steady state, except for the interest rate and the inflation rate, which are both shown as percentage point deviations from the risky steady state and expressed in annualized terms.⁴

3.1 A Change in the Inflation Target

The inflation target in central banking models—particularly within the New Keynesian framework—significantly influences the conduct and effectiveness of monetary policy. By raising its inflation target, a central bank signals a greater tolerance for higher inflation. In the short run, this policy can stimulate output by lowering the real interest rate and encouraging spending; however, in the long run, it can lead to higher actual inflation and its associated economic costs.

The optimal level of the inflation target remains a subject of ongoing research and debate. It depends on various factors, including the structure of the economy, the nature of shocks, and the effectiveness of unconventional monetary policy tools. Recent literature has focused on the potential benefits and costs of raising the inflation target above the commonly observed 2% level, particularly in light of persistent low-interest-rate environments and the heightened risk of hitting the ZLB. Nonetheless, any adjustment to the inflation target must be carefully considered, weighing the potential macroeconomic benefits against the risks to central bank credibility.

In our first bToTEM experiment, we analyze an anticipated change in the inflation target within the Taylor rule (eq:23). Specifically, we assume the central bank announces in advance that it will increase the inflation target $\bar{\pi}_t$, and that economic agents consider this announcement to be fully credible. Why is this a relevant policy experiment for the Canadian economy? During the Great Recession of 2007–2009, many central banks lowered their nominal policy rates to levels they considered to be

4. By risky steady state, we mean the state to which a stochastic economy converges when the exogenous shocks are set to zero.




Fig. 5. A gradual increase in the inflation target.

the ZLB (see Witmer and Yang 2015 for an updated estimate of Canada's ZLB). An extensive literature argues that the inflation target is an effective policy instrument for managing ZLB episodes. For example, Summers (1991) and Fischer (1996) suggest maintaining an inflation target as high as 2% or 3% if the economy hits the ZLB, while Krugman (1998) proposes a 4% inflation target for the Japanese economy to combat persistent deflation. Similarly, Blanchard, Dell'Ariccia, and Mauro (2010), Williams (2009), and Ball (2013) argue that a higher inflation target could have prevented the interest rate from reaching the ZLB.

Canada adopted an inflation-targeting framework in 1991, aiming for a gradual decline in the inflation target to 2% by the end of 1995. Since then, the inflation target has been consistently maintained at 2%. This inflation-targeting framework is reviewed and renewed every 5 years. Notably, during the most recent review in 2021, the Bank of Canada opted to keep the target at the 2% level; this renewal covers the period through the end of 2026. This environment creates two types of possible anticipation effects. First, if the Bank of Canada had decided to change the target in 2021, the policy implementation lag would have generated anticipation effects. Second, even though the inflation target remained unchanged in 2021, anticipation effects were nonetheless present, as there was a nonzero probability of a target adjustment given Canada's proximity to the ZLB at the time and the serious policy discussions surrounding this possibility.

Figure 5 displays the dynamics of the main model's variables. We present the results for the method that finds a perturbation solution obtained around a deterministic steady state (labeled as Method 2 in online Appendix B; our sensitivity results for other methods predict similar patterns of behavior).

We consider two scenarios: First, the central bank makes an announcement that it will gradually increase the inflation target π from 2% to 3% over a period of eight quarters, starting at $t = 1$. Second, the same gradual increase occurs but begins at $t = 5$ (i.e., in 1 year). In both cases, the inflation target remains at the new, higher level once the adjustment is complete.

When the inflation-target change begins at $t = 1$, inflation mirrors the trajectory of the target. This behavior occurs because, in our experiment, we assume full credibility of the inflation-target policy. Inflation repeats the pattern of the inflation target because the agents determining the behavior of inflation are primarily nonoptimizers who index their prices to the inflation target. Consequently, the nominal interest rate gradually increases over the first 15 periods by 1%, and it stays at the new level forever (see Figure 5; note that the real neutral rate remains unchanged). Following the announcement, output, investment, and commodity exports experience an immediate increase. During the transition, the economy undergoes investment- and export-driven growth, with output peaking at a 0.2% increase. However, output begins to decline toward its original level after 1 year, resulting in only a temporary expansionary effect on the economy from the higher inflation target.⁵

When the inflation-target change is delayed by 1 year, the variables exhibit qualitatively similar behavior to the immediate-change case. However, one notable difference is that most variables in the figure experience larger increases at their peaks. In particular, output increases to about 0.3%, with the exceptions of the exchange rate and noncommodity exports. This result suggests that announcing such a policy in advance benefits the economy, as output sees greater gains during the transition period.

The larger jumps in variables such as output, consumption, investment, and capital are entirely driven by anticipatory effects. Agents expect the real interest rate to be lower in the near future and choose to accumulate more capital ahead of the more favorable economic environment. This forward-looking behavior positively impacts the economy's output even before the inflation-target change takes effect.

Table 1 contains the mean and maximum residuals in the model's equations used for computing the corresponding variables in the table. As shown, the maximum residuals range between $10^{-3.14}$ and $10^{-5.85}$, that is, between 0.07% and 0.0001%, which are exceptionally low.⁶ For the remaining experiments, the residuals in the

5. Garin, Lester, and Sims (2018) analyze the impact of inflation-target shocks on output within a canonical three-equation New Keynesian model. Their findings indicate that the output response varies significantly depending on the persistence of the shock (see their Figure 2). For sufficiently persistent shocks, the effect of the inflation target on output is quantitatively comparable to the results in our bToTEM model. However, because their shock is persistent but stationary, it differs from the nonstationary structures studied here. A closer alignment is found in Uribe (2022), who examines a permanent, nonstationary increase in the inflation target within an estimated small New Keynesian model and identifies short-run expansionary effects on output. Importantly, neither of these studies explores the implications of postponing the inflation-target shock, which remains a distinctive feature of our EFP analysis.

6. In LMM (2020), we reported larger residuals when the economy was hit by a large negative demand shock and the ELB was reached.

TABLE 1
RESIDUALS IN THE MODEL'S EQUATIONS ON THE SIMULATED PATH, $\log_{10}$ UNITS

	R_t	π_t	Y_t	C_t	I_t	X_t^{nc}	X_t^{com}	M_t	L_t	K_t
Average	-4.75	-5.25	-4.09	-4.47	-5.11	-4.10	-3.17	-5.09	-4.42	-5.91
Maximum	-4.40	-5.19	-4.04	-4.40	-5.04	-4.09	-3.14	-4.99	-4.38	-5.85

NOTE: $R_t, \pi_t, Y_t, C_t, I_t, X_t^{nc}, X_t^{com}, M_t, L_t, K_t$ are the nominal interest rate, inflation, output, consumption, investment, noncommodity export, commodity export, imports, labor, and capital, respectively.




Fig. 6. A gradual increase in the inflation target (50% probability).
equations are of similar size, meaning our solutions are highly accurate (to save on space, these residuals are not reported).

In the second experiment, we model a probabilistic setting in which agents rationally expect that the inflation target may take one of two possible levels with given probabilities. Specifically, we assume a 50% probability that, starting from $t = 5$, the inflation target π_t will gradually increase from 2% to 3% over eight quarters; otherwise, the inflation target remains unchanged. Our computational method is straightforward to adapt to more complex anticipation scenarios such as the one considered in this experiment. To compute policies in period $t = 4$, we explicitly use the Dynare macro language to set the period 4 expectations equal to the weighted sums of the two possible outcomes in period $t = 5$.

This experiment is illustrated in Figure 6, showing two alternative transition paths diverging from period 5 onward, one per each scenario—that is, with and without an increase in $\bar{\pi}_t$. Similar to the previous experiment, inflation closely follows the behavior of the inflation target: it gradually rises to the new steady-state level. Starting from the risky steady state at $t = 1$, all variables exhibit mild increases, driven by anticipa-

tory effects on the part of economic agents. Once it is revealed whether the inflation target will change, variables rapidly revert to the original steady state if the target remains unchanged. Conversely, if the target increases, variables display a more pronounced hump-shaped adjustment before stabilizing at the new steady state. In online Appendix B, we extend this experiment by varying the probability of switching to a higher inflation target at $t = 5$, analyzing probabilities of 25% and 75% (instead of 50%). The results indicate that in the absence of an inflation-target change, the transition back to the old steady state is significantly faster when the *ex ante* probability of switching is 25% compared to 75%.

3.2 Monetary Policy Normalization

Monetary policy normalization refers to the central bank's actions to raise nominal interest rates back to a standard rule after a period of being at the ZLB with an accommodative monetary policy. It has several effects within the New Keynesian model. First, monetary policy normalization leads to a decrease in output and a decrease in inflation in the short run because higher interest rates dampen aggregate demand by making borrowing more expensive and saving more attractive. Second, it can affect financial markets by leading to higher bond yields and potentially lower equity prices, as higher interest rates make bonds more attractive relative to other assets. Overall, the impact of monetary policy normalization within the New Keynesian model depends on various factors, including the initial state of the economy, the pace and communication of the normalization process, and the credibility of the central bank. If the economy is close to the ZLB, the effectiveness of monetary policy normalization may be limited. In such cases, the central bank may need to rely on unconventional monetary policy tools, such as forward guidance and quantitative easing, to stimulate the economy. By announcing their intentions to eventually raise rates, central banks can guide expectations and influence economic behavior even before the actual policy change occurs.

In Canada, as in many other countries, the nominal interest rate reached the ZLB during the Great Recession of 2007–2009. Consequently, central banks were unable to rely on Taylor rules to guide their monetary policy and instead turned to forward guidance—an unconventional monetary policy consisting of announcing future interest-rate changes. As emphasized in the literature, the primary mechanism through which forward guidance influences the economy is the central bank's communication of its intended policy-rate path. Eggertsson and Woodford (2003) demonstrate that a central bank's promises to maintain low interest rates for extended periods help alleviate the negative consequences of a binding ZLB. As agents expect future interest rates to be lower than in the absence of forward guidance, they increase current investment and consumption, thereby stimulating the economy in the present. Campbell et al. (2012) identify two types of forward guidance: Odyssean and Delphic. Odyssean forward guidance refers to explicit commitments by the central bank to future policy actions, ensuring that its intentions are unequivocally clear. In contrast, Delphic forward guidance involves the central bank providing a




Fig. 7. Monetary policy normalization.

forecasted policy-rate path based on its information regarding the economic shocks at play. Unlike the Odyssean approach, Delphic forward guidance reveals intended future policy actions without committing to them, leaving agents to interpret the guidance.⁷ In this paper, we focus on Odyssean forward guidance.

Specifically, we assume that the central bank uses forward guidance to convey a policy change when lifting off from an ELB on its policy rates. In particular, we assume that the economy is initially at the ELB, and at $t = 1$, the central bank announces its intention to maintain the interest rate at that level for T periods before reverting to the standard Taylor rule (23). To model the central bank's policy during the ELB periods, we assume that the nominal interest rate is given by

$$R_t = R^{elb},$$

where R^{elb} is the ELB level of the nominal interest rate. When the interest-rate policy is normalized after T periods, the Taylor rule's coefficients return to their normal values, and the policy is described by the rule (23).

Figure 7 presents the results for this experiment when the solutions are approximated around a deterministic steady state. The change in the interest-rate rule announced at $t = 1$ is anticipated by agents. We compare three cases, depending on whether the interest-rate policy returns to normal (i) in one quarter ($T = 1$), (ii) in 1

7. Marinkov (2020) argues that during the ZLB period, agents may misjudge a central bank's reaction function and bias their expectations. In such cases, the central bank may want to use forward guidance as a tool to correct agents' beliefs.

year ($T = 4$), or (iii) in 2 years ($T = 8$). In all cases, the initial interest rate is below its risky steady state; however, it eventually returns to the steady state.

When the policy is announced, the exchange rate, inflation, and real variables jump above the steady state. Local currency depreciation makes domestic exports more competitive, leading to an increase in exports of both commodities and noncommodity goods. Domestic firms benefit from increased sales, which leads to immediate increases in output, labor, investment, and capital. On the other hand, as households work more, they demand more imported goods, causing imports to rise as well. Evidently, the output increase is largest when the announced policy is maintained for the longest horizon of eight quarters. The differences in output dynamics across the considered cases are only present during the transition, not in the initial period—in all three cases, the initial output jump is roughly of equal size. The independence of the initial reaction in output from the horizon of the forward guidance policy is related to the well-known “forward guidance puzzle” in the literature. Even though there is no such dependence in the figure, we still observe that the policy horizon matters for the total effect on output: the economy reacts more strongly if the policy change is postponed further into the future.

Our experiment contributes to the ongoing discussion regarding central banks’ communication strategies. Following the Great Recession of 2007–2009, as economic conditions improved, a critical policy question emerged: how and when should monetary policy normalization occur? Here, normalization refers to transitioning back to a standard Taylor rule; see Yellen (2015). Specifically, the following questions arose post-crisis:

- (1) Should the central bank normalize policy immediately or later?
- (2) Should normalization be implemented gradually or all at once?
- (3) Should the regime shift be announced in advance?
- (4) Should policy normalization be time-dependent or state-dependent?

These questions are difficult to address within the conventional stationary New Keynesian framework because monetary policy normalization inherently involves nonstationary changes. However, the techniques developed in this paper allow us to study these questions effectively. In our above experiment, we compare the economy’s behavior under policies with differing horizons for returning to standard rules, directly addressing questions (1) and (3). Our findings indicate that announcing a future lift-off in advance provides measurable benefits, and we quantify these gains for various durations of forward guidance. Questions (2) and (4) can similarly be explored using our methods; we leave them for future research.

An effective commitment to keeping the interest rate at the ELB implies that the rate should be maintained at this low level for longer than a standard Taylor rule would prescribe. Notably, during the COVID-19 pandemic, interest rates reached the ZLB across a number of developed economies. Taylor (2021) argues that the Taylor rules considered by the Fed in the February 2021 Monetary Policy Report implied that the federal funds rate should have been higher than the actual zero level and that the Fed “should now engage in a strategy or rule in which people and markets

understand that it would raise the policy interest rate if economic growth increases and inflation rises as they are now forecast to do.” Our analysis, however, emphasizes the paramount importance of commitment to the announced policy, upon which the desired monetary expansion ultimately hinges.

3.3 *Switching to a More Aggressive Taylor Rule*

A more aggressive Taylor rule implies that the central bank responds more strongly to changes in inflation and output, yielding several key implications for central banking models. First, a more aggressive Taylor rule (with a coefficient on inflation greater than 1) is often crucial for achieving a unique and stable equilibrium. It anchors inflation expectations and mitigates the risks of inflationary and deflationary spirals. Second, while a more aggressive Taylor rule promotes stability, it can also lead to increased volatility in output and interest rates in the short run. Because the central bank responds more forcefully to economic fluctuations, these interventions can inadvertently amplify the initial shocks. Third, a credible commitment to a more aggressive Taylor rule can enhance the central bank’s ability to manage expectations, leading to lower and more stable inflation in the long run, even with a less volatile policy path. Fourth, a stronger response to inflation can complicate the identification of policy effects, as it becomes harder to disentangle the impact of the rule from other factors influencing inflation dynamics. Overall, a more aggressive Taylor rule serves as a powerful tool for achieving price stability and output stabilization; however, its implementation requires careful consideration of the associated trade-offs and challenges.

In our bToTEM experiment, we consider a one-time change in the sensitivity of the policy rate to inflation and the output gap within the Taylor rule (23), represented by ρ_π and ρ_Y , respectively.⁸ This experiment differs from our previous analyses, in which the anticipated changes in the model’s parameters were implemented gradually. Figure 8 illustrates the economy’s responses to twofold increases in either ρ_π or ρ_Y , or both, relative to the benchmark parameterization. It is worth noting that this change in the coefficient values is quite large relative to what a central bank would typically consider. The switch to a more aggressive Taylor rule is anticipated at $t = 1$ but occurs at $t = 2$. Consequently, all model variables exhibit immediate anticipatory effects in response to this scheduled policy adjustment.

Both policies—a higher ρ_π and a higher ρ_Y —are inflationary. However, a twofold increase in the sensitivity to inflation ρ_π is more effective at expanding the economy. Output, consumption, investment, capital, and labor are visibly higher both at their peaks and in the new steady state compared to the old steady state. Conversely, commodity production slightly declines, likely due to reduced commodity exports. In contrast, a twofold increase in the sensitivity to the output gap yields more modest effects. When there is a stronger response to both inflation and the output gap simultaneously, the quantitative magnitude of the effects falls roughly between the

8. For example, Taylor (1999) argues that the Taylor rule with $\rho_\pi = 0.5$ and $\rho_Y = 1$ is more reasonable than the one advocated in Taylor (1993), where $\rho_\pi = 0.5$ and $\rho_Y = 0.5$.




Fig. 8. A switch to more aggressive Taylor rules.

two individual cases. This outcome reflects the trade-off between inflation and the output gap within the policy rule: a stronger response to the output gap offsets some of the effects of a stronger response to inflation.

We also observe differences in the risky steady-state interest rate. When the Taylor rule features a stronger response to the inflation gap, the risky steady-state interest rate is lower. Conversely, it is higher when the response to the output gap is stronger. The intuition behind this result is as follows: when the response to the inflation gap becomes stronger, a smaller increase in the interest rate is required to achieve equivalent inflation stabilization. Similarly, when the response to the output gap becomes stronger, a larger increase in the interest rate is needed to achieve the same degree of inflation stabilization. Overall, these total effects are not quantitatively significant. Switching to a significantly more aggressive Taylor rule has only minor effects on the economy's behavior in the absence of shocks.

3.4 Switching from Inflation Targeting to PLT

Switching from inflation targeting to PLT has several implications for central banking models. First, PLT has the potential to reduce the volatility of output and inflation through automatic, self-correcting changes in inflation expectations, which could be particularly beneficial during periods of economic stress. Second, PLT can be especially useful at the ZLB, as it allows for a catch-up period of above-target inflation following a period of below-target inflation, helping to stimulate the economy when conventional monetary policy tools are constrained. However, PLT requires a strong focus on managing public expectations, as its effectiveness hinges on agents' forward-looking behavior and their understanding of the policy; this ne-

cessitates greater transparency and communication from central banks. Furthermore, the success of PLT depends heavily on its credibility. If the public doubts the central bank's commitment to the price-level path, the policy may not be as effective in stabilizing the economy. Overall, while switching from inflation targeting to PLT offers significant potential benefits—such as enhanced economic stability and better management of ZLB scenarios—the transition poses challenges. Establishing public understanding and building credibility for the new framework are critical and may require significant time and effort.

In the literature, the seminal paper by Svensson (1999) argues that PLT is a “free lunch” in the sense that it positively affects the short-run trade-off between inflation and output variability (namely, it reduces inflation variability without an increase in output variability). Hatcher and Minford (2016) and Ambler (2009) provide comprehensive surveys on this topic, further highlighting the advantages of PLT.

Bernanke (2017) proposes using a temporary price-level target when short-term interest rates are at or near the ZLB. At the ZLB, policymakers are often constrained in their ability to provide adequate stimulus, leading to below-target inflation. A temporary price-level target addresses this issue by compensating for periods of below-target inflation with a subsequent temporary surge in inflation. The Bank of Canada has seriously considered the use of PLT; see Kahn (2009) and Bank of Canada (2011).

We first consider a bToTEM scenario, in which central bank switches from the standard Taylor rule (23) targeting inflation to the one targeting a price-level gap,

$$R_t = \rho_r R_{t-1} + (1 - \rho_r) \{ \bar{R} + \rho_\pi (\log P_t - \log \bar{P}_t) + \rho_Y (\log Y_t - \log \bar{Y}_t) \} + \eta_t^r, \quad (37)$$

where P_t is the actual price level, and $\bar{P}_t$ is the target price level that grows at the rate of inflation target $\bar{P}_t = \bar{P}_{t-1} \bar{\pi}_t$. Therefore, PLT does not imply that policymakers pursue a constant price level, but rather that they set a target for the price level that rises over time.

An inflation-targeting central bank tolerates temporary deviations from the inflation target, as long as inflation eventually returns to the target level, effectively “letting bygones be bygones.” In contrast, a PLT central bank actively works to reverse temporary deviations from the inflation target whenever they occur. For example, if inflation falls below the target, the central bank increases inflation to bring the price level back to its predetermined path. Under inflation targeting, an inflation shock permanently shifts the price path to a different level, while under PLT, any movement in inflation above target is matched with an equal and opposite movement in inflation below target, ensuring the economy follows a predetermined price path. As a result, inflation targeting introduces significant uncertainty about the future price level because the central bank only gradually returns inflation to the target level without correcting for past misses. In contrast, PLT provides much greater certainty about future price levels, as it commits to correcting deviations and adhering to the established price path, even with a positive average inflation rate.

In Figure 9, we present the results of two policy experiments where the policy change becomes effective either immediately $t = 1$ or 1 year after being announced




Fig. 9. A switch to price-level targeting.

at $t = 5$. The new interest-rate rule is associated with higher steady-state levels for all the model's variables in the figure. Therefore, switching to PLT has expansionary effects on the economy. Moreover, for all variables except the real interest rate, the immediately implemented policy yields larger benefits than the policy announced 1 year in advance. That is, if the central bank postpones implementing the switch, the transition effects are somewhat smaller.

In the next experiment, we shock the economy, creating a large output gap. In particular, we consider a permanent negative demand shock—a decrease in foreign demand, modeled as a negative innovation in the random-walk process for this shock. (A version of this experiment with a permanent decrease in productivity is presented in online Appendix B.) In response, the central bank can either continue using a policy rule with inflation targeting or switch to PLT, which, as previously demonstrated, leads to a higher steady-state output. As a result, switching to PLT can be viewed as an attempt to revive the economy.

Figure 10 presents the results of this experiment where the switch to PLT is either implemented immediately or delayed for 1 year (but still announced today). As shown, there is minimal difference between the immediate and delayed changes in monetary policy for variables such as commodities and commodity exports. For all other variables, the immediate policy change has a larger impact than the delayed policy change, consistent with the previous experiment in Figure 9.

Therefore, the anticipation effects work to soften the impact of the negative demand shock, with impulse responses falling between the cases of no change and immediate change. In particular, the dynamics of the nominal interest rate are smoother in the case of anticipated policy, which leads to smoother behavior in the remaining variables. The initial impact on the economy is significant: output and labor fall by




Fig. 10. A negative foreign demand shock and a switch to price-level targeting.

1% and 1.5% in the three scenarios considered. With no switch in monetary policy, output recovers slightly, but its new steady state remains below the old steady state. With the policy switch, output returns to nearly the same level as before the shock, and consumption is even higher. Notably, each considered policy scenario involves the central bank tightening monetary policy, even during the economic downturn.

3.5 Switching from Inflation Targeting to AIT

Switching from inflation targeting to AIT introduces several important changes to central banking models:

First, AIT allows inflation to temporarily deviate from the target, provided that over a set period, the average inflation rate returns to the target. This flexibility is useful for responding to economic shocks and mitigating short-term fluctuations, which can help stabilize the economy in turbulent times.

Second, by committing to a long-run average inflation target, policymakers can enhance the credibility of their commitment to price stability and anchor long-run inflation expectations. This can lead to more stable and predictable inflation dynamics.

Third, the higher flexibility inherent in AIT enables policymakers to better stabilize output and employment in response to economic shocks. This can yield better economic outcomes—especially in times of crisis—by avoiding excessive policy tightening in the short run and focusing on long-term stability.

Fourth, AIT is particularly beneficial in the context of the ZLB. Allowing for temporary overshooting of the inflation target can help mitigate the constraints imposed by the ZLB, making monetary policy more effective during economic downturns.

Fifth, in certain scenarios, AIT shares properties with PLT, such as reducing inflation and output volatility. This suggests that AIT may offer some of the benefits of PLT—such as enhancing economic stability—without the same degree of implementation challenges.

Finally, the success of AIT depends on the central bank's ability to effectively communicate the policy framework to the public and maintain its credibility. Misunderstandings or doubts about the central bank's commitment to the long-term average target could undermine the policy's effectiveness.

Overall, AIT provides central banks with more tools to manage economic fluctuations and maintain stability, while also offering a degree of flexibility not present in traditional inflation targeting. However, like any policy shift, its effectiveness depends on clear communication and the public's trust in the central bank's commitment.

The U.S. Federal Reserve's transition from inflation targeting to AIT was formally announced by Fed Chair Jerome Powell on August 27, 2020; see Powell (2020). However, as Richard Clarida noted during his presentation at the Hoover Economic Policy Working Group on January 13, 2021, there was evidence that the Fed would introduce that framework as early as a month before the formal announcement. This anticipation led to substantial price adjustments in the U.S. economy even before the official policy shift.

In our bToTEM experiment, we consider a switch from the inflation-targeting Taylor rule (23) to an alternative rule that incorporates an average of past inflation (including current actual inflation). The modified rule is as follows:

$$R_t = \rho_r R_{t-1} + (1 - \rho_r) \left[\bar{R} + \rho_\pi \left(\frac{1}{M+1} \sum_{j=0}^M \pi_{t-j} - \bar{\pi}_t \right) + \rho_Y (\log Y_t - \log \bar{Y}_t) \right] + \eta'_t. \quad (38)$$

The policy of AIT shares several characteristics with PLT. As suggested by the previous literature, AIT is a middle ground between PLT and inflation targeting; see Nessén and Vestin (2005). Under AIT, a central bank reacts to a deviation of the average of current and past inflation from the inflation target. For example, if the inflation target is 2%, the averaging window is 3 years, and after consistently achieving 2% inflation in the past, inflation deviates to 3% in the most recent year, the central bank will aim to achieve a policy-induced inflation of 1% in the following year. As a result, inflation will oscillate around the average inflation target, ensuring the target is met on average. The price level will stay close to its trend, even though the level will sometimes deviate from the fixed trend.

Figure 11 displays the outcomes of two scenarios in our experiment: one where the switch happens immediately and another where it is implemented with a lag of 1 year after it was announced. Amano et al. (2020) study the optimal degree of history dependence under AIT in the context of the standard New Keynesian model that includes the ELB. Their analysis concludes that the optimal length of the averaging window, denoted as M , ranges from 2 to 8. In this experiment, we adopt $M = 8$, the maximum value identified in their study. This policy change exhibits very modest anticipation effects on the economy in the absence of any shocks. Notably, when the policy becomes effective immediately, variables such as output, labor, imports, and




Fig. 11. A switch to average inflation targeting.

noncommodity exports show stronger responses compared to a delayed implementation. This suggests that reacting to average inflation rather than contemporaneous inflation smooths out the dynamics to a new steady state. Therefore, we would not expect the economy to experience any drastic changes during the course of the transition to AIT.

4. MARKOV VERSUS NON-MARKOV NEWS SHOCKS

Other methods in the literature that focus on changes in the economic environment include Markov regime-switching models (e.g., Davig and Leeper 2007) and models with Markov news shocks (e.g., Schmitt-Grohé and Uribe 2012); see also Barro and King (1984), Beaudry and Portier (2006), 2007, and Jaimovich and Rebelo (2009). However, there is a fundamental conceptual difference between this literature and the EFP framework. The existing literature considers recurrent changes driven by a Markov process and constructs one time-invariant decision rule, whereas we consider a sequence (path) of non-Markov events and construct a sequence (path) of time-varying decision rules. This distinction is crucial, as the stationary solutions generated by Markov-based methods may differ significantly from the nonstationary solutions derived from non-Markov analysis. The quantitative differences can be substantial. A single decision rule designed for a Markov process cannot adequately capture the variety of non-Markov scenarios that may occur along a historical path. This highlights the EFP framework's ability to better address the complexities and dynamics of nonstationary, time-dependent economic environments.

To illustrate this difference, we compare the solutions of Schmitt-Grohé and Uribe (2012) with those derived using the EFP framework in the context of the experiment in Section 4.3, where the central bank switches to a more aggressive Taylor rule. Specifically, we assume that the sensitivity to inflation ρ_π in the Taylor rule (23) is doubled relative to its benchmark value. The policy change is announced at $t = 1$ and implemented at $t = 2$. Our framework with non-Markov shocks provides a natural approach to modeling this scenario. Using our method, we construct a stationary solution for period $t = 2$, and we find a solution for period $t = 1$ that matches a given terminal condition (decision rule) constructed for period $t = 2$. This sequential, time-varying approach allows us to capture the transitional dynamics and anticipatory effects that are inherently nonstationary, contrasting sharply with the stationary solutions typically generated by Markov-based methods.

Schmitt-Grohé and Uribe (2012) do not specify how their perturbation method can be applied to analyze non-Markov anticipated shocks. To adapt their method to our experiment, we explored two approaches. First, we consider a unit-root process for $\rho_{\pi,t}$, that is, $\rho_{\pi,t} = \rho_{\pi,t-1} + \varepsilon_{t-1}$, in which initially $\rho_{\pi,0} = \rho_\pi$. In this specification, the shock innovation ε_t captures news that becomes known at period t and has a direct impact at $t + 1$. In our experiment, $\varepsilon_1 = \rho_\pi$ at period $t = 1$, and in all other periods, the shock innovation is zero. This implies that $\rho_{\pi,1} = \rho_\pi$ and $\rho_{\pi,t} = 2\rho_\pi$ for all $t \geq 2$. Second, we consider a temporary news shock, that is, $\rho_{\pi,t} = \rho_\pi + \varepsilon_{t-1}$. In that case, we obtain $\rho_{\pi,2} = 2\rho_\pi$ in period $t = 2$ and $\rho_{\pi,t} = \rho_\pi$ in all other periods. These adaptations attempt to incorporate non-Markov features into the Schmitt-Grohé and Uribe (2012) framework. However, they rely on stationary decision rules and do not fully capture the time-varying nature of non-Markov scenarios as modeled in our EFP framework.

Figure 12 compares our second-order perturbation EFP solution with two solutions produced by Schmitt-Grohé and Uribe (2012) method. The volatility of the news shock is assumed to be zero, so the initial risky steady state is the same for all three solutions. The following observations are in order. First, the EFP solution with non-Markov shocks is situated between the two Markov news-shock solutions. Second, both the EFP solution and the permanent Markov news-shock solution converge to new risky steady states, although these steady states differ. In contrast, the temporary Markov news-shock solution reverts to the old steady state, consistent with the transient nature of the shock. Third, across all three cases, the variables behave in a qualitatively similar manner. A more aggressive central bank leads to an increase in inflation and a decrease in nominal and real interest rates, which raises output, investment, capital, and imports. Fourth, the gap between our solution and the permanent news-shock solution varies with the initial condition. Sensitivity experiments (not reported here) reveal that the gap is smaller if we start below the steady state. This occurs because anticipation effects are intertwined with upward-sloping transition dynamics. Finally, the difference between the two solutions with Markov news shocks and our second-order perturbation solution stems from differences in the slopes of the decision rules. The EFP approach accounts for non-Markov features, resulting in unique slopes that better capture the dynamics of the system. If first-order perturba-




Fig. 12. A switch to a more aggressive Taylor rule at $t = 2$ announced at $t = 1$.

tion methods were used, the economy would remain at the deterministic steady state in the two Markov news-shock solutions, but not in our solution.

Furthermore, we observe that the permanent Markov news-shock model predicts dramatically larger effects associated with the switch to a more aggressive Taylor rule than the EFP method, with the exception of consumption. This discrepancy is evident both in the anticipation effects and in the differences between steady states. For example, the anticipation effects on investment are five times larger at their peak for the Markov news shocks than for our perturbation solutions. We conclude that the two approaches can lead to qualitatively different results. Specifically, the Markov news-shock approach tends to overstate the importance of a given anticipated event. This overestimation arises because the unit-root process for $\rho_{\pi,t}$ in the Markov framework implies that once a news shock occurs, its effects will persist forever.

There is a simple intuition for why our non-Markov solutions differ from those produced for Markov news shocks. In our framework, the agent's decision rule represents the best response to a specific sequence of news shocks. In contrast, in Schmitt-Grohé and Uribe (2012) framework, it is the best response to a given Markov stationary process. With a Markov process, the response to news is determined not only by the news itself but also by the properties of the Markov process assumed when constructing the solution. This dependency on the underlying Markov process introduces dynamics that are absent in our non-Markov framework. For the same reason, regime-switching Markov models are not well-suited for analyzing non-Markov news shocks, as they inherently rely on the Markov assumption, which fails to capture the path-dependent nature of non-Markov scenarios.

5. CONCLUSION

The literature recognizes that the private sector's expectations are critically important for policy outcomes, and central banks frequently use open-mouth policies to anchor these expectations. However, little work has been done on evaluating the effects of open-mouth policies within a Dynamic Stochastic General Equilibrium (DSGE) framework. This paper fills this gap. We find that anticipation effects are strongest for time-dependent economic policies, such as policy-rate normalization in the aftermath of a ZLB crisis and a gradual change in the inflation target level. Other time-dependent policy changes, such as a switch to a more aggressive policy-rate rule or a shift to price-level or AIT, lead to more modest anticipation effects.

The proposed EFP methodology is not limited to central banking models. Many economic policies are announced well ahead of being implemented. For instance, changes to taxes, tariffs, minimum wages, pension reforms, and Social Security are frequently signed into law long before they are put into practice. Other notable examples include announcements regarding a new member state's accession to the European Union (EU), a member state's exit from the EU (e.g., Brexit), or the outcome of a presidential election before the newly elected administration takes power. Our perturbation-based framework for solving, calibrating, simulating, and parameter estimation provides a simple and tractable way to analyze the non-Markov transitions associated with such policy changes. Ultimately, the EFP analysis makes it possible to construct a model-consistent historical path of a real-world economy.

LITERATURE CITED

- Amano, Robert, Stefano Gnocchi, Sylvain Leduc, and Joel Wagner. (2020) "Average is Good Enough: Average-Inflation Targeting and the ELB." Bank of Canada Staff Working Paper 2020–31.
- Ambler, Steve. (2009) "Price-Level Targeting and Stabilization Policy: A Survey." *Journal of Economic Surveys*, 23, 974–97.
- Ball, Laurence. (2013) "The Case for Four Percent Inflation." *Central Bank Review*, 13, 17–31.
- Bank of Canada. (2011) "Renewal of the Inflation-Control Target. Background information – November 2011."
- Barro, Robert, and Robert King. (1984) "Time Separable Preferences and Intertemporal Substitution Models of Business Cycles." *Quarterly Journal of Economics*, 99, 817–39.
- Beaudry, Paul, and Franck Portier. (2006) "News, Stock Prices, and Economic Fluctuations." *American Economic Review*, 96, 1293–307.
- Beaudry, Paul, and Franck Portier. (2007) "When can Changes in Expectations Cause Business Cycle Fluctuations in Neo-Classical Settings?" *Journal of Economic Theory*, 135, 458–77.
- Bernanke, Ben. (2017) "Temporary Price-Level Targeting: An Alternative Framework for Monetary Policy." Brookings, October 12, 18–26.

- Blanchard, Olivier, Giovanni Dell’Ariccia, and Paolo Mauro. (2010) “Rethinking Macroeconomic Policy.” International Monetary Fund Staff Position Note, SPN/10/03.
- Brock, William. (1971) “Sensitivity of Optimal Growth Paths with Respect to a Change in Target Stocks.” In *Contributions to the Von Neumann Growth Model*, edited by Gerhart Bruckmann and Wilhelm Weber, pp. 73–89. New York: Springer Verlag.
- Brock, William, and Leonard Mirman. (1972) “Optimal Economic Growth and Uncertainty: The Discounted Case.” *Journal of Economic Theory*, 4, 479–513.
- Campbell, Jeffrey R., Charles L. Evans, Jonas D. Fisher, Alejandro Justiniano, Charles W. Calomiris, and Michael Woodford. (2012) “Macroeconomic Effects of Federal Reserve Forward Guidance [with comments and discussion].” *Brookings Papers on Economic Activity*, 1–80.
- Corrigan, Paul, Hélène Desgagnés, José Dorich, Vadym Lepetyuk, Wataru Miyamoto, and Yang Zhang. (2021) “ToTEM III: The Bank of Canada’s Main DSGE Model for Projection and Policy Analysis.” Bank of Canada Technical Report No. 119.
- Davig, Troy, and Eric Leeper. (2007) “Generalizing the Taylor Principle.” *American Economic Review*, 97, 607–35.
- Dorich, José, Michael Johnston, Rhys Mendes, Stephen Murchison, and Ying Zhang. (2013) “ToTEM II: An Updated Version of the Bank of Canada’s Quarterly Projection Model.” Bank of Canada Technical Report No. 100.
- Eggertsson, Gauti, and Michael Woodford. (2003) “The Zero Bound on Interest Rates and Optimal Monetary Policy.” *Brookings Papers on Economic Activity*, 2003, 139–211.
- Fair, Ray, and John Taylor. (1983) “Solution and Maximum Likelihood Estimation of Dynamic Nonlinear Rational Expectations Models.” *Econometrica*, 51, 1169–85.
- Fischer, Stanley. (1996) “Why are Central Banks Pursuing Long-Run Price Stability?” *Economic Policy Symposium, Jackson Hole, Proceedings*, Federal Reserve Bank of Kansas City, 7–34.
- Gali, Jordi, and Mark Gertler. (1999) “Inflation Dynamics: A Structural Econometric Analysis.” *Journal of Monetary Economics*, 44, 195–222.
- Garin, Julio, Robert Lester, and Eric Sims. (2018) “Raise Rates to Raise Inflation? Neo-Fisherianism in the New Keynesian Model.” *Journal of Money, Credit and Banking*, 50, 243–59.
- Hatcher, Michael, and Patrick Minford. (2016) “Stabilization Policy, Rational Expectations and Price-Level Versus Inflation Targeting: A Survey.” *Journal of Economic Surveys*, 30, 327–55.
- Jaimovich, Nir, and Sergio Rebelo. (2009) “Can News About the Future Drive the Business Cycle?” *American Economic Review*, 99, 1097–118.
- Joshi, Sumit. (1997) “Turnpike Theorems in Nonconvex Nonstationary Environments.” *International Economic Review*, 38, 245–48.
- Kahn, George. (2009) “Beyond Inflation Targeting: Should Central Banks Target the Price Level?” *Economic Review*, Federal Reserve Bank of Kansas City, Third Quarter, 35–65.
- King, Robert, Charles Plosser, and Sergio Rebelo. (1988) “Production, Growth and Business Cycles.” *Journal of Monetary Economics*, 21, 195–232.
- Krugman, Paul. (1998) “It’s Baaack: Japan’s Slump and the Return of the Liquidity Trap.” *Brookings Papers on Economic Activity*, 2, 137–87.

- Lepetyuk, Vadym, Lilia Maliar, and Serguei Maliar. (2020) "When the U.S. Catches a Cold, Canada Sneezes: A Lower-Bound Tale Told By Deep Learning." *Journal of Economic Dynamics and Control*, 117, 103926.
- Majumdar, Mukul, and Itzhak Zilcha. (1987) "Optimal Growth in a Stochastic Environment: Some Sensitivity and Turnpike Results." *Journal of Economic Theory*, 43, 116–33.
- Maliar, Lilia, Serguei Maliar, John B. Taylor, and Inna Tsener. (2020) "A Tractable Framework for Analyzing a Class of Nonstationary Markov Models." *Quantitative Economics*, 11, 1289–323.
- Marinkov, Vesna. (2020) "Communication at the Zero Lower Bound: The Case for Forward Guidance." Manuscript.
- Mitra, Tapan, and Yaw Nyarko. (1991) "On the Existence of Optimal Processes in Nonstationary Environments." *Journal of Economics*, 53, 245–70.
- Nakamura, Emi, and Jón Steinsson. (2018) "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect." *Quarterly Journal of Economics*, 133, 1283–330.
- Nessén, Marianne, and David Vestin. (2005) "Average Inflation Targeting." *Journal of Money, Credit and Banking*, 37, 837–63.
- Powell, Jerome. (2020) "New Economic Challenges and the Fed's Monetary Policy Review." Economic Policy Symposium "Navigating the Decade Ahead: Implications for Monetary Policy."
- Schmitt-Grohé, Stephanie, and Martín Uribe. (2003) "Closing Small Open Economy Models." *Journal of International Economics*, 61, 163–85.
- Schmitt-Grohé, Stephanie, and Martín Uribe. (2004) "Solving Dynamic General Equilibrium Models Using a Second-Order Approximation to the Policy Function." *Journal of Economic Dynamics and Control*, 28, 755–75.
- Schmitt-Grohé, Stephanie, and Martín Uribe. (2012) "What's News in Business Cycles." *Econometrica*, 80, 2733–64.
- Summers, Lawrence. (1991) "Panel Discussion: Price Stability: How Should Long-Term Monetary Policy Be Determined?" *Journal of Money, Credit and Banking*, 23(3 Part 2), 625–31.
- Svensson, Lars. (1999) "Price-Level Targeting Versus Inflation Targeting: A Free Lunch?" *Journal of Money, Credit and Banking*, 31, 277–95.
- Taylor, John B. (1993) "Discretion Versus Policy Rules in Practice." *Carnegie-Rochester Conference Series on Public Policy*, 39, 195–214.
- Taylor, John B. (1999) "A Historical Analysis of Monetary Policy Rules." In *Monetary Policy Rules*, edited by John B. Taylor, pp. 655–79. Chicago, IL: University of Chicago Press.
- Taylor, John B. (2021) "The Optimal Reentry to a Monetary Policy Strategy." Manuscript.
- Uribe, Martín. (2022) "The Neo-Fisher Effect: Econometric Evidence from Empirical and Optimizing Models." *American Economic Journal: Macroeconomics*, 14, 133–62.
- Williams, John. (2009) "Heeding Daedalus: Optimal Inflation and the Zero Lower Bound." *Brookings Papers on Economic Activity*, 2, 1–37.
- Yellen, Janet. (2015) "Normalizing Monetary Policy: Prospects and Perspectives." Board of Governors of the Federal Reserve System. Speech.
- Yi, Kei-Mu, and Jing Zhang. (2017) "Understanding Global Trends in Long-Run Real Interest Rates." *Economic Perspectives*, 41.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of the article.

APPENDIX A: EFP FOR THE GROWTH MODEL

Table A1: Difference of a simulated solution path from the balanced growth path in $\log_{10}$ units

Figure A1: Comparison of the nonstationary EFP solutions computed by Method 5 and the balanced growth solution

Table A2: Comparison of the EFP to the other methods

APPENDIX B: SENSITIVITY EXPERIMENTS

APPENDIX C: CALIBRATION