Andreasen, Martin Møller ; Kronborg, Anders Farver

Article

The extended perturbation method: With applications to the New Keynesian model and the zero lower bound

Quantitative Economics

Provided in Cooperation with:
The Econometric Society

Suggested Citation: Andreasen, Martin Møller ; Kronborg, Anders Farver (2022) : The extended perturbation method: With applications to the New Keynesian model and the zero lower bound, Quantitative Economics, ISSN 1759-7331, The Econometric Society, New Haven, CT, Vol. 13, Iss. 3, pp. 1171-1202,
https://doi.org/10.3982/QE1102

This Version is available at:
https://hdl.handle.net/10419/296298

Quantitative Economics 13 (2022), 1171–1202 1759-7331/20221171

The extended perturbation method: With applications to the New Keynesian model and the zero lower bound

Martin M. Andreasen
Department of Economics and Business Economics Aarhus University, CREATES, and Danish Finance Institute

Anders F. Kronborg
Danish Research Institute for Economic Analysis and Modelling

We introduce the extended perturbation method, which improves the accuracy of standard perturbation by reducing approximation errors under certainty equivalence. For the New Keynesian model with Calvo pricing, extended perturbation is more accurate than standard perturbation, which implies explosive dynamics because it omits the upper bound on inflation implied by this model. In contrast, extended perturbation enforces this bound and generates stable dynamics. We also show that extended perturbation can accurately solve a New Keynesian model that enforces the zero lower bound for the monetary policy rate by considering a smooth nonlinear modification of the standard Taylor rule.

Keywords. Stability of perturbation approximations, upper bound on inflation, zero lower bound on policy rate.

JEL classification. C62, C68, E30.

1. Introduction

The solution to dynamic stochastic general equilibrium (DSGE) models is frequently approximated by the perturbation method to obtain higher-order Taylor series expansions of the policy function. The popularity of this approximation is mainly explained by its ability to (i) preserve nonlinearities in the model, (ii) improve parameter identification compared to a linearized solution, and (iii) capture effects of uncertainty to explore determinants of risk premia and implications of uncertainty shocks (see An and Schorfheide (2007), Kim and Ruge-Murcia (2009), Fernández-Villaverde, Guerrón-Quintana, Rubio-Ramírez, and Uribe (2011), Rudebusch and Swanson (2012), among others).

Martin M. Andreasen: mandreasen@econ.au.dk
Anders F. Kronborg: ank@dreamgruppen.dk
We thank Rhys Bidder, Lawrence Christiano, Vasco Curdia, Wouter den Haan, Jens Iversen, Anders B. Kock, Ales Marsal, Eric Swanson, Allan Sørensen, Konstantinos Theodoridis, and Joris de Wind for useful comments and discussions. M. Andreasen greatly acknowledges access to computer facilities at the Centre for Scientific Computing in Aarhus (CSCAA) and funding from the Independent Research Fund Denmark, Project 7024-00020B, and support from the Danish Finance Institute (DFI) and the Center for Research in Econometric Analysis of Time Series (CREATES).

 https://doi.org/10.3982/QE1102

Despite the widespread use of second- and third-order perturbation approximations, it is well known that they often generate explosive dynamics even when the corresponding linearized solution is stable. The perturbation approximation may also struggle to preserve key properties of the true solution such as monotonicity and convexity, as emphasized by Haan and De Wind (2012). These findings suggest that the second- and third-order perturbation approximations currently applied in the literature may not always be sufficiently accurate. A tractable alternative that preserves the stability of the true solution, but not necessarily its monotonicity and convexity, is to apply a pruning scheme, as proposed by Kim, Kim, Schaumburg, and Sims (2008) for models approximated to second order and extended to higher order by Haan and De Wind (2012), Andreasen, Fernandez-Villaverde, and Rubio-Ramirez (2018), and Lombardo and Uhlig (2018).

The contribution of the present paper is to improve the accuracy of the perturbation approximation by combining it with the extended path of Fair and Taylor (1983). This is done by decomposing the policy function into (i) a deterministic component under certainty equivalence and (ii) a stochastic component containing the effects of uncertainty. The perturbation method is currently applied to approximate *both* parts of the policy function. However, a relatively low Taylor series expansion of the certainty equivalent component (e.g., to second or third order) may generate unnecessary approximation errors, given that this part of the policy function can be computed very accurately by the extended path at a small computational cost. Based on this observation, we propose to compute the certainty equivalent component of the policy function by the extended path, whereas the stochastic part of the policy function remains approximated by standard perturbation. We name this combined solution method extended perturbation, which improves the accuracy of standard perturbation by reducing approximation errors under certainty equivalence.

The approximation order in extended perturbation is determined by the order of the polynomial used for the stochastic part of the policy function. For a first- and second-order approximation, we show that extended perturbation gives a stable solution, provided the model is stable under certainty equivalence. This result does not generalize beyond second order due to the local nature of the uncertainty correction from standard perturbation. We therefore present a test for stability, and a simple procedure to ensure stability if the approximation does not pass this stability test.

Using a standard New Keynesian model with Calvo pricing, we show that extended perturbation is more accurate than standard perturbation when using approximations up to fifth order. We also find that extended perturbation is stable even when standard perturbation explodes. To understand this difference, we use our stability test to locate critical state configurations that lead to explosive dynamics. The analysis shows that standard perturbation is unable to account for an upper bound on inflation with Calvo pricing, and this leads to an explosive price-inflation spiral in the approximation. This upper bound on inflation has been largely ignored in the literature, but extended perturbation properly accounts for this bound, and hence generates a stable approximation.

In an empirical application, we enforce the zero lower bound (ZLB) for the monetary policy rate in the New Keynesian model by proposing a smooth nonlinear modification of the existing Taylor rule. Higher-order perturbation methods are well suited

for solving this model, because the proposed Taylor rule does not introduce a kink, as implied by the widely used shadow rate specification (see Gust, Herbst, Lopez-Salido, and Smith (2017, Arouba, Cuba-Borda, and Schorfheide (2018), among many others). We then show that this modified version of the New Keynesian model can generate long stays at the ZLB and is easy to solve using extended perturbation, which provides a very accurate solution.

The paper proceeds as follows. Section 2 presents the extended perturbation method, and Section 3 studies its stability. Section 4 analyzes accuracy and stability of this approximation for a standard New Keynesian model. The computational issues related to extended perturbation are discussed in Section 5, and Section 6 presents our ZLB application. Section 7 concludes.[1]

2. The extended perturbation method

We start by introducing the considered class of DSGE models in Section 2.1. The extended perturbation method is presented in Section 2.2, while Sections 2.3 and 2.4 provide an analytical and numerical example, respectively. We finally compare extended perturbation to other approximation methods in Section 2.5.

2.1 *DSGE models*

We consider the broad class of DSGE models, which can be expressed as

$$\mathbb{E}_t\big[\mathbf{f}(\mathbf{x}_t, \mathbf{x}_{t+1}, \mathbf{y}_t, \mathbf{y}_{t+1})\big] = \mathbf{0}, \tag{1}$$

where $\mathbb{E}_t$ denotes the conditional expectation given information available in period t. The state vector $\mathbf{x}_t$ with dimension $n_x \times 1$ belongs to the Borel set $\mathcal{X}_x \subseteq \mathbb{R}^{n_x}$. The control variables are stored in $\mathbf{y}_t$ with dimension $n_y \times 1$ and $\mathbf{y}_t \in \mathcal{X}_y$, where $\mathcal{X}_y$ refers to the Borel subset of $\mathbb{R}^{n_y}$. We further let $n_x + n_y = n$. The function $\mathbf{f}$ maps elements from $\mathcal{X}_x \times \mathcal{X}_x \times \mathcal{X}_y \times \mathcal{X}_y$ into $\mathbb{R}^n$, and we assume that this mapping is at least m times differentiable, where m will be used below to indicate the approximation order of the model.

We further let $\mathbf{x}_t \equiv [\mathbf{x}_{1,t}' \ \mathbf{x}_{2,t}']'$, where $\mathbf{x}_{1,t}$ contains the endogenous state variables and $\mathbf{x}_{2,t}$ denotes the exogenous states. The dimensions of these vectors are $n_{x_1} \times 1$ and $n_{x_2} \times 1$, respectively, with $n_{x_1} + n_{x_2} = n_x$. The dynamics of $\mathbf{x}_{2,t}$ are given by

$$\mathbf{x}_{2,t+1} = \mathbf{\Gamma}(\mathbf{x}_{2,t}) + \sigma\bar{\boldsymbol{\eta}}\boldsymbol{\epsilon}_{t+1}, \tag{2}$$

where the function $\mathbf{\Gamma}$ maps elements from the Borel set $\mathcal{X}_{x_2} \subseteq \mathbb{R}^{n_x}$ into $\mathcal{X}_{x_2}$ and is required to be at least m times differentiable. The innovations $\boldsymbol{\epsilon}_{t+1}$ belong to the Borel set $\mathcal{X}_{\boldsymbol{\epsilon}} \subseteq \mathbb{R}^{n_x}$ and has dimension $n_{\boldsymbol{\epsilon}} \times 1$. These innovations are assumed to be independent and identically distributed with zero mean and a unit covariance matrix, that is, $\boldsymbol{\epsilon}_{t+1} \sim \mathcal{IID}(\mathbf{0}, \mathbf{I})$. We also require that each element of $\boldsymbol{\epsilon}_{t+1}$ has a finite mth moment to compute the standard perturbation approximation up to mth order. The function $\mathbf{\Gamma}$

[1]We also provide a MATLAB package that implements the extended perturbation approximation up to fifth order.

must be specified such that it generates a stable process for $\mathbf{x}_{2,t}$.[2] In linear systems, this corresponds to requiring that all eigenvalues of the Jacobian $\partial \boldsymbol{\Gamma} / \partial \boldsymbol{x}_{2,t}'$ lie inside the unit circle. For nonlinear systems, $\boldsymbol{\Gamma}$ must satisfy the general stability condition for nonlinear first-order Markov processes provided below in Section 3.1.

We focus on models with a unique and globally stable solution that has one steady state. This implies that we do not consider models with multiple steady states, local indeterminancy as in Lubik and Schorfheide (2004), and models with limit cycles around an unstable steady state as in Beaudry, Galizia, and Portier (2020), which can be solved by various modifications of standard perturbation.

The exact solution to the considered class of DSGE models may then be expressed as

$$\mathbf{y}_t = \mathbf{g}(\mathbf{x}_t, \sigma), \tag{3}$$

$$\mathbf{x}_{t+1} = \mathbf{h}(\mathbf{x}_t, \sigma) + \sigma \boldsymbol{\eta} \boldsymbol{\epsilon}_{t+1}, \tag{4}$$

$$\boldsymbol{\eta} \equiv \begin{bmatrix} \mathbf{0}_{n_{x_1} \times n_\epsilon} \\ \bar{\boldsymbol{\eta}} \end{bmatrix}.$$

The assumption that innovations enter linearly in (2) and (4) is without loss of generality, because the state vector may be extended to account for nonlinearities between $\mathbf{x}_t$ and $\boldsymbol{\epsilon}_{t+1}$ as shown by Andreasen (2012). The perturbation parameter $\sigma \geq 0$ scales the square root of the covariance matrix for the innovations $\boldsymbol{\eta}$ with dimension $n_x \times n_\epsilon$ and enables us to capture effects of uncertainty in the policy functions. In particular, when $\sigma = 0$ we get the deterministic model under certainty equivalence, that is,

$$\begin{aligned} \mathbf{g}^{\text{CE}}(\mathbf{x}_t) &\equiv \mathbf{g}(\mathbf{x}_t, \sigma = 0), \\ \mathbf{h}^{\text{CE}}(\mathbf{x}_t) &\equiv \mathbf{h}(\mathbf{x}_t, \sigma = 0), \end{aligned} \tag{5}$$

whereas the model with uncertainty is obtained by letting $\sigma = 1$. Unfortunately, the policy functions $\mathbf{g}$ and $\mathbf{h}$ in (3) and (4) are generally unknown and must be approximated.

2.2 *The extended perturbation method*

This paper builds on the observation that the policy functions can be decomposed into two parts: (i) a certainty equivalent solution and (ii) a stochastic component, which we define as the difference between the true solution and the certainty equivalent solution. That is,

$$\begin{aligned} \mathbf{g}^{\text{stoch}}(\mathbf{x}_t, \sigma) &\equiv \mathbf{g}(\mathbf{x}_t, \sigma) - \mathbf{g}^{\text{CE}}(\mathbf{x}_t), \\ \mathbf{h}^{\text{stoch}}(\mathbf{x}_t, \sigma) &\equiv \mathbf{h}(\mathbf{x}_t, \sigma) - \mathbf{h}^{\text{CE}}(\mathbf{x}_t), \end{aligned} \tag{6}$$

[2]This implies that trends may only be included in the considered class of DSGE models if a given model after rescaling has an equivalent representation without trending variables. A similar requirement is needed to apply the standard perturbation method (see, for instance, King and Rebelo (1999)).

where $\mathbf{g}^{\text{stoch}}$ and $\mathbf{h}^{\text{stoch}}$ denote the stochastic part of the policy function, as indicated by the superscript. Inserting (6) into (3) and (4), the exact solution may be expressed as

$$\mathbf{y}_t = \mathbf{g}^{\text{CE}}(\mathbf{x}_t) + \mathbf{g}^{\text{stoch}}(\mathbf{x}_t, \sigma),$$
$$\mathbf{x}_{t+1} = \mathbf{h}^{\text{CE}}(\mathbf{x}_t) + \mathbf{h}^{\text{stoch}}(\mathbf{x}_t, \sigma) + \sigma\boldsymbol{\eta}\boldsymbol{\epsilon}_{t+1}.$$

Following the work of Judd and Guu (1997), the perturbation method is usually applied to approximate both ($\mathbf{g}^{\text{CE}}$, $\mathbf{h}^{\text{CE}}$) and ($\mathbf{g}^{\text{stoch}}$, $\mathbf{h}^{\text{stoch}}$) at the steady state (ss), that is, at $\mathbf{x}_{\text{ss}} = \mathbf{x}_{t+1} = \mathbf{x}_t$ and $\sigma = 0$. However, a relatively low Taylor series expansion of $\mathbf{g}^{\text{CE}}$ and $\mathbf{h}^{\text{CE}}$ to second or third order may generate unnecessary approximation errors, given that $\mathbf{g}^{\text{CE}}$ and $\mathbf{h}^{\text{CE}}$ can be computed very accurately by the extended path at a small computational cost. This is done by setting all innovations to their expected value of zero and imposing a terminal condition $\mathbf{y}_{t+N}$ at some finite horizon N for the system in (1). For a given state $\mathbf{x}_t$, the extended path then finds $\{\mathbf{x}_{t+i}^{\text{EP}}\}_{i=1}^{N}$ and $\{\mathbf{y}_{t+i}^{\text{EP}}\}_{i=0}^{N-1}$ that satisfy the equilibrium conditions in (1) in the absence of uncertainty for the next N periods, implying that $\mathbf{g}^{\text{CE}}(\mathbf{x}_t) = \mathbf{y}_t^{\text{EP}}$ and $\mathbf{h}^{\text{CE}}(\mathbf{x}_t) = \mathbf{x}_{t+1}^{\text{EP}}$.[3]

Thus, we suggest to solve for $\mathbf{g}^{\text{CE}}$ and $\mathbf{h}^{\text{CE}}$ by the extended path, but continue to approximate the stochastic part of the policy function $\mathbf{g}^{\text{stoch}}$ and $\mathbf{h}^{\text{stoch}}$ by standard perturbation. This alternative might appear to be a somewhat ad hoc combination of a global solution method (the extended path) and a local method (the perturbation approximation), but the procedure does have a proper theoretical foundation. To realize this, we first show in the Appendix that the uncertainty adjustments implied by the perturbation approximation remain valid, as they also match the curvature of the true solution at the steady state when $\mathbf{g}^{\text{CE}}$ and $\mathbf{h}^{\text{CE}}$ are approximated by the extended path instead of Taylor series expansions. Second, an infinite Taylor series expansion of the certainty equivalent solution to the $\mathbf{g}$-function at the steady state implies

$$\mathbf{g}^{\text{CE}}(\mathbf{x}_t) = \sum_{k=0}^{\infty} \frac{\mathbf{g}(\mathbf{x}_{\text{ss}}, 0)_{\mathbf{x}^k}}{k!}(\mathbf{x}_t - \mathbf{x}_{\text{ss}})^{\otimes k},$$

where $\mathbf{g}(\mathbf{x}_{\text{ss}}, 0)_{\mathbf{x}^k}$ with dimension $n_y \times (n_x)^k$ denotes partial derivatives of $\mathbf{g}$ at the steady state taken k times with respect to $\mathbf{x}_t$ and $(\mathbf{x}_t - \mathbf{x}_{\text{ss}})^{\otimes k} = (\mathbf{x}_t - \mathbf{x}_{\text{ss}})^{\otimes(k-1)} \otimes (\mathbf{x}_t - \mathbf{x}_{\text{ss}})$. A similar expression obviously applies for the $\mathbf{h}$-function. This shows that our combined solution procedure may be interpreted as adding higher-order certainty equivalent terms to a standard perturbation approximation. For instance, in the case of a third-order approximation ($m = 3$), we add all the higher-order terms $\sum_{k=m+1}^{\infty} \frac{g(x_{\text{ss}},0)_{\mathbf{x}^k}}{k!}(x_t - x_{\text{ss}})^k$. Thus, our procedure may be considered as a tractable way to "extend" the standard perturbation approximation by additional higher-order terms that are otherwise not possible to include due to computational constraints. Based on this observation, we name our combined solution procedure the extended perturbation method, which improves the accuracy of standard perturbation by reducing approximation errors in the certainty equivalent part of the policy function.

[3] Further details on the extended path are provided below in Section 5.1.

The approximation order for extended perturbation is determined by the order of the Taylor series expansion used to approximate $\mathbf{g}^{\text{stoch}}$ and $\mathbf{h}^{\text{stoch}}$. We illustrate this by presenting approximations up to fifth order for the $\mathbf{g}$-function, while similar expressions apply for the $\mathbf{h}$-function. Here and throughout, we use subscripts to denote partial derivatives of a given function when evaluated at the steady state. A first-order approximation simply reproduces the certainty equivalent solution by the extended path. This is because the partial derivatives $\mathbf{g}_{\mathbf{x}}^{\text{stoch}}$ and $\mathbf{g}_{\sigma}^{\text{stoch}}$ are zero at the steady state (see Schmitt-Grohé and Uribe (2004)).

At second order, we have

$$\mathbf{y}_t = \mathbf{g}^{\text{CE}}(\mathbf{x}_t) + \frac{1}{2}\mathbf{g}_{\sigma\sigma}\sigma^2, \tag{7}$$

where $\mathbf{g}_{\sigma\sigma}$ captures a constant uncertainty correction due to variance risk. The cross-derivative $\mathbf{g}_{\sigma\mathbf{x}}$ is zero as shown in Schmitt-Grohé and Uribe (2004), implying that the uncertainty correction is not state dependent at second order.

The third-order approximation reads

$$\mathbf{y}_t = \mathbf{g}^{\text{CE}}(\mathbf{x}_t) + \frac{1}{2}\mathbf{g}_{\sigma\sigma}\sigma^2 + \frac{3}{6}\mathbf{g}_{\sigma\sigma\mathbf{x}}\sigma^2(\mathbf{x}_t - \mathbf{x}_{\text{ss}}) + \frac{1}{6}\mathbf{g}_{\sigma\sigma\sigma}\sigma^3, \tag{8}$$

as $\mathbf{g}_{\sigma\mathbf{xx}} = \mathbf{0}$ (see Andreasen (2012)). Thus, $\mathbf{g}_{\sigma\sigma\mathbf{x}}$ captures a linear state-dependent uncertainty correction due to variance risk, while $\mathbf{g}_{\sigma\sigma\sigma}$ corrects the level of the approximation for skewness risk; for instance induced by rare disasters as in Rietz (1988).

The fourth-order approximation is given by

$$\begin{aligned}\mathbf{y}_t = \mathbf{g}^{\text{CE}}(\mathbf{x}_t) &+ \frac{1}{2}\mathbf{g}_{\sigma\sigma}\sigma^2 + \frac{1}{6}\mathbf{g}_{\sigma\sigma\sigma}\sigma^3 + \frac{1}{24}\mathbf{g}_{\sigma\sigma\sigma\sigma}\sigma^4 \\ &+ \left[\frac{3}{6}\mathbf{g}_{\sigma\sigma\mathbf{x}}\sigma^2 + \frac{4}{24}\mathbf{g}_{\sigma\sigma\sigma\mathbf{x}}\sigma^3\right](\mathbf{x}_t - \mathbf{x}_{\text{ss}}) \\ &+ \frac{6}{24}\mathbf{g}_{\sigma\sigma\mathbf{xx}}\sigma^2(\mathbf{x}_t - \mathbf{x}_{\text{ss}})^{\otimes 2},\end{aligned} \tag{9}$$

as $\mathbf{g}_{\sigma\mathbf{xxx}} = \mathbf{0}$. The term $\mathbf{g}_{\sigma\sigma\mathbf{xx}}$ captures a quadratic correction for variance risk, while $\mathbf{g}_{\sigma\sigma\sigma\mathbf{x}}$ represents a state-dependent linear correction for skewness risk. Finally, $\mathbf{g}_{\sigma\sigma\sigma\sigma}$ corrects the level of the approximation for kurtosis risk.

The fifth-order approximation takes the form

$$\begin{aligned}\mathbf{y}_t = \mathbf{g}^{\text{CE}}(\mathbf{x}_t) &+ \frac{1}{2}\mathbf{g}_{\sigma\sigma}\sigma^2 + \frac{1}{6}\mathbf{g}_{\sigma\sigma\sigma}\sigma^3 + \frac{1}{24}\mathbf{g}_{\sigma\sigma\sigma\sigma}\sigma^4 + \frac{1}{120}\mathbf{g}_{\sigma\sigma\sigma\sigma\sigma}\sigma^5 \\ &+ \left[\frac{3}{6}\mathbf{g}_{\sigma\sigma\mathbf{x}}\sigma^2 + \frac{4}{24}\mathbf{g}_{\sigma\sigma\sigma\mathbf{x}}\sigma^3 + \frac{5}{120}\mathbf{g}_{\sigma\sigma\sigma\sigma\mathbf{x}}\sigma^4\right](\mathbf{x}_t - \mathbf{x}_{\text{ss}}) \\ &+ \left[\frac{6}{24}\mathbf{g}_{\sigma\sigma\mathbf{xx}}\sigma^2 + \frac{10}{120}\mathbf{g}_{\sigma\sigma\sigma\mathbf{xx}}\sigma^3\right](\mathbf{x}_t - \mathbf{x}_{\text{ss}})^{\otimes 2} \\ &+ \frac{10}{120}\mathbf{g}_{\sigma\sigma\mathbf{xxx}}\sigma^2(\mathbf{x}_t - \mathbf{x}_{\text{ss}})^{\otimes 3},\end{aligned} \tag{10}$$

as $\mathbf{g}_{\sigma\mathbf{xxxx}} = \mathbf{0}$. The term $\mathbf{g}_{\sigma\sigma\sigma\sigma\mathbf{x}}$ includes a linear correction for kurtosis risk, while $\mathbf{g}_{\sigma\sigma\sigma\mathbf{xx}}$ provides a quadratic adjustment for skewness risk. Variance risk is further accounted for by the term $\mathbf{g}_{\sigma\sigma\mathbf{xxx}}$, and $\mathbf{g}_{\sigma\sigma\sigma\sigma\sigma}$ modifies the level of the approximation for risks related to the fifth moment.

The presence of the perturbation parameter σ in (7) to (10) is included to emphasize that the two components of the policy function $\mathbf{g}^{\mathrm{CE}}(\mathbf{x}_t)$ and $\mathbf{g}^{\mathrm{stoch}}(\mathbf{x}_t, \sigma)$ capture different and mutually exclusive aspects of the solution. The certainty equivalence component $\mathbf{g}^{\mathrm{CE}}(\mathbf{x}_t)$ is computed in the absence of uncertainty with $\sigma = 0$, which imply that all uncertainty corrections in (7) to (10) are zero; however, the presence of uncertainty with $\sigma = 1$ does not affect $\mathbf{g}^{\mathrm{CE}}(\mathbf{x}_t)$ by construction, but it allows the uncertainty corrections in (7) to (10) to affect the solution.

Thus, it is straightforward to form the mth order approximation by the extended perturbation method. The required steps are:

Step 1: Use standard perturbation to obtain all derivatives of $\mathbf{g}^{\mathrm{stoch}}(\mathbf{x}_t, \sigma)$ and $\mathbf{h}^{\mathrm{stoch}}(\mathbf{x}_t, \sigma)$ at the steady state to order m.

Step 2: In any time period, use the extended path to compute $\mathbf{g}^{\mathrm{CE}}(\mathbf{x}_t)$ and $\mathbf{h}^{\mathrm{CE}}(\mathbf{x}_t)$ and approximate $\mathbf{g}^{\mathrm{stoch}}(\mathbf{x}_t, \sigma)$ and $\mathbf{h}^{\mathrm{stoch}}(\mathbf{x}_t, \sigma)$ by $\hat{\boldsymbol{g}}^{\mathrm{stoch}}(\mathbf{x}_t, \sigma)$ and $\hat{\boldsymbol{h}}^{\mathrm{stoch}}(\mathbf{x}_t, \sigma)$, respectively, using the derivatives from standard perturbation.

2.3 *An analytical example*

We illustrate our approximation method by applying it to the asset pricing model of Lucas (1978). A representative agent is here maximizing $E_0[\sum_{t=0}^{\infty} \beta^t \frac{c_t^\theta}{\theta}]$ with respect to consumption c_t subject to $p_t e_{t+1} + c_t = (p_t + d_t)e_t$, where e_t is the amount of the asset held in period t at price p_t. Dividends d_t are given by $d_{t+1} = d_t \exp\{x_{t+1}\}$, where $x_{t+1} = (1-\rho)x_{\mathrm{ss}} + \rho x_t + \sigma\eta\epsilon_{t+1}$ and $\epsilon_{t+1} \sim \mathcal{NID}(0,1)$. Denoting the price-dividend ratio by $y_t = p_t/d_t$, Burnside (1998) shows that the solution is

$$y_t \equiv g(x_t, \sigma) = \sum_{i=1}^{\infty} \beta^i \exp\bigl\{a_i + b_i(x_t - x_{\mathrm{ss}})\bigr\},$$

where $a_i = \theta x_{\mathrm{ss}} i + \frac{\theta^2\sigma^2\eta^2}{2(1-\rho)^2}[i - \frac{2\rho(1-\rho^i)}{1-\rho} + \frac{\rho^2(1-\rho^{2i})}{1-\rho^2}]$ and $b_i = \frac{\theta\rho(1-\rho^i)}{1-\rho}$. The certainty equivalent solution is obtained by letting $\sigma = 0$, implying that $g^{\mathrm{CE}}(x_t) = \sum_{i=1}^{\infty} \beta^i \exp\{\theta x_{\mathrm{ss}} i + b_i(x_t - x_{\mathrm{ss}})\}$. Using the decomposition in (6), we have

$$g(x_t, \sigma) = g^{\mathrm{CE}}(x_t) + \underbrace{\sum_{i=1}^{\infty} \beta^i \bigl(\exp\bigl\{a_i + b_i(x_t - x_{\mathrm{ss}})\bigr\} - \exp\bigl\{\theta x_{\mathrm{ss}} i + b_i(x_t - x_{\mathrm{ss}})\bigr\}\bigr)}_{g^{\mathrm{stoch}}(x_t, \sigma)}.$$

This analytic expression for $g^{\mathrm{stoch}}(x_t, \sigma)$ shows that the stochastic part of the policy function varies nonlinearly with x_t. It is also easy to verify the results for the partial derivatives of $\mathbf{g}^{\mathrm{stoch}}(\mathbf{x}_t, \sigma)$ stated above. For instance, we clearly have $\partial g^{\mathrm{stoch}}(x_t, \sigma)/\partial x_t|_{\mathrm{ss}} = 0$ and $\partial g^{\mathrm{stoch}}(x_t, \sigma)/\partial\sigma|_{\mathrm{ss}} = 0$ at the steady state, showing that a first-order approximation to $g^{\mathrm{stoch}}(x_t, \sigma)$ does not include any correction for uncertainty.

2.4 *A computational example*

The key insight exploited by extended perturbation is that the extended path can be a more efficient way to capture the information in higher-order derivatives of $\mathbf{g}^{\text{CE}}(\mathbf{x}_t)$ and $\mathbf{h}^{\text{CE}}(\mathbf{x}_t)$ than the standard perturbation method. We illustrate this point by considering the multicountry real business cycle model in Juillard and Villemot (2011) with k countries. This model has $2k+1$ controls variables and $2k$ state variables, meaning that the size and computational cost of the model is determined by k. We refer to our Appendix in the Online Supplementary Material for further details about the model and the considered calibration. Panel $\mathcal{A}$ in Table 1 shows that the number of seconds to obtain the standard perturbation solution increases rapidly with the approximation order for models with more than ten states. For a fifth-order approximation, it takes 79.0 seconds with $k=8$, 297.8 seconds (≈ 5 min.) with $k=10$ and 1322.6 seconds (≈ 22 min.) with $k=12$. Panel $\mathcal{B}$ shows that simply simulating a sample path of 1000 observations may also be costly when using a high approximation order. At fifth order, it takes 17.9 seconds with $k=8$, 68.2 seconds with $k=10$, and 198.8 seconds with $k=12$. This rapid increase in the computing time is mainly because evaluating the fifth-order terms related to $(\mathbf{x}_t-\mathbf{x}_{\text{ss}})^{\otimes 5}$ becomes increasingly expensive with many states. The execution time of simulating 1000 observations using the extended path increases much more slowly with the size of the model and takes only 55.6 seconds with $k=10$ and 74.6 seconds with $k=12$, which in both cases are *faster* than the time used by the corresponding fifth-order perturbation approximation. Panel $\mathcal{C}$ reports the root mean squared errors (RMSE^{CE}) for the unit-free Euler-equation errors along the same simulated sample path. To focus on the certainty equivalence part of the solution, we evaluate these errors for a deterministic version of the model and transform them by the 10-base logarithm ($\log_{10}$) to facilitate the interpretation. The accuracy of standard perturbation improves monotonically with the order of the approximation and is about 10^{-5} at fifth order, but it remains less accurate when compared to the extended path with a RMSE^{CE} of about 10^{-15}.

Accordingly, the time needed to compute very high perturbation approximations (say, to fifth order) are substantial for models with more than 10 state variables. In addition, one should also account for sizable costs of running simulations at very high orders of approximation, which are comparable to the time used by a benchmark version of the extended path. In Section 5.1, we discuss several refinements of this benchmark version of the extended path to substantially reduce its execution time.

2.5 *Comparing extended perturbation to the literature*

The extended perturbation method is related to the stochastic version of the extended path in Fair and Taylor (1983), where the extended path is modified to account for uncertainty. A recent application of this modification is considered in Adjemian and Juillard (2013), where they account for uncertainty about the innovations in the first $S \geq 1$ periods when running the extended path, after which all innovations are set to zero. For these first S periods, the model equations implied by (1) are evaluated by numerical integration using Gaussian quadratures, and this makes the method of Adjemian and Juillard (2013) computationally demanding because S must be substantially larger than one

TABLE 1. A multicountry model: computational costs and accuracy. This table reports the number of seconds for computing the standard perturbation solution (panel $\mathcal{A}$) and for simulating a sample path of $T = 1000$ observations (panel $\mathcal{B}$). The accuracy along the simulated sample path from the extended path is reported in panel $\mathcal{C}$ by the log10-transformed root mean squared unit-free Euler-equation errors for a deterministic version of the model. The standard perturbation solution is obtained using the codes of Levintal (2017). For the extended path, the starting values and the controls at the end horizon ($N = 50$) are determined from a first-order approximation. All computations are carried out on a standard desktop with Windows 10 using an Intel(R) Core(TM) i7-7600U CPU @ 280 GHz processor with 16 GB memory.

	$k = 4$	$k = 6$	$k = 8$	$k = 10$	$k = 12$
Panel $\mathcal{A}$: Seconds for perturbation solution					
Perturbation: 1st-order	0.00	0.01	0.02	0.08	0.19
Perturbation: 2nd-order	0.00	0.01	0.02	0.08	0.19
Perturbation: 3rd-order	0.00	0.01	0.03	0.11	3.28
Perturbation: 4th-order	0.35	1.21	4.07	9.92	24.3
Perturbation: 5th-order	1.69	16.3	79.0	297.8	1322.6
Panel $\mathcal{B}$: Seconds for simulating $T = 1000$					
Perturbation: 1st-order	0.00	0.04	0.12	0.44	0.93
Perturbation: 2nd-order	0.00	0.05	0.13	0.44	1.02
Perturbation: 3rd-order	0.00	0.10	0.18	0.58	1.80
Perturbation: 4th-order	0.00	0.43	1.10	3.98	9.31
Perturbation: 5th-order	0.00	4.89	17.9	68.2	198.8
Extended path	14.0	21.0	31.6	55.6	74.6
Panel $\mathcal{C}$: Accuracy by $\log_{10}$ (RMSE^{CE})					
Perturbation: 1st-order	−2.8	−2.8	−2.9	−3.1	−3.1
Perturbation: 2nd-order	−3.8	−3.9	−4.0	−4.2	−4.2
Perturbation: 3rd-order	−4.3	−4.2	−4.4	−4.6	−4.5
Perturbation: 4th-order	−5.0	−4.9	−5.2	−5.4	−5.3
Perturbation: 5th-order	−5.4	−5.2	−5.6	−5.8	−5.7
Extended path	−14.9	−14.9	−15.2	−15.2	−15.2

to obtain a reasonable degree of accuracy (e.g., about 30). To reduce the execution time of their approximation, Adjemian and Juillard (2013) suggest to select a fairly low value of S and instead apply the second-order uncertainty correction from standard perturbation to risk correct all endogenous variables beyond S periods when computing the extended path. That is, $\{\mathbf{y}_{t+i} + \frac{1}{2}\mathbf{g}_{\sigma\sigma}\}_{i=S+1}^{N}$ and $\{\mathbf{x}_{1,t+i} + \frac{1}{2}\mathbf{h}_{1,\sigma\sigma}\}_{i=S+1}^{N}$ are used instead of $\{\mathbf{y}_{t+i}\}_{i=S+1}^{N}$ and $\{\mathbf{x}_{1,t+i}\}_{i=S+1}^{N}$ when running the extended path. In contrast, extended perturbation does not apply the uncertainty correction from standard perturbation when running the extended path, but instead add these risk corrections to the certainty equivalent solution implied by the extended path.

Another way to combine the extended path with the perturbation method is presented in Ajevskis (2017). He augments the certainty equivalent solution by a pointwise second-order accurate estimate of $\mathbf{g}^{\text{stoch}}$ and $\mathbf{h}^{\text{stoch}}$. Extended perturbation instead modifies the solution from the extended path by adding the perturbation approximation of $\mathbf{g}^{\text{stoch}}$ and $\mathbf{h}^{\text{stoch}}$ to this solution. Thus, the loadings for the uncertainty correction in extended perturbation are only computed at the steady state, whereas the uncertainty

correction in Ajevskis (2017) is limited to second order in the current implementation and must be recomputed at every value of $\mathbf{x}_t$, making this method computationally more demanding than extended perturbation in long simulations.

3. A stability test for smooth approximations

An important feature of any approximation method is to preserve stability of the exact solution to the considered class of DSGE models. This stability requirement is easy to examine in a linear approximation by inspecting the eigenvalues of the state transition function. However, a similar stability condition is rarely discussed in relation to nonlinear approximations. We therefore start in Section 3.1 by presenting a stability condition for nonlinear Markov systems, which should be of general interest for researchers working with nonlinear approximations. This stability condition is then applied in Section 3.2 to analyze the stability of extended perturbation. Section 3.3 introduces a numerical test for stability, which is relevant for extended perturbation beyond second order where stability cannot be guaranteed. A simple procedure to ensure stability of extended perturbation beyond second order is finally provided in Section 3.4.

3.1 *Stability of a nonlinear Markov system*

This section presents a stability condition for the system in (3) and (4). It is here sufficient to focus on the process for $\mathbf{x}_t$, given that the control variables are a smooth function of the states. We apply the stability condition in Pötscher and Prucha (1997) for the first-order nonlinear Markov system in (4). To present this condition, iterate (4) forward by k periods to obtain

$$\mathbf{x}_{t+k} = \mathbf{h}^{(k)}(\mathbf{x}_t, \boldsymbol{\epsilon}_{t+1}, \boldsymbol{\epsilon}_{t+2}, \dots \boldsymbol{\epsilon}_{t+k-1}, \sigma) + \sigma\boldsymbol{\eta}\boldsymbol{\epsilon}_{t+k},$$

where $\mathbf{h}^{(2)}(\mathbf{x}_t, \boldsymbol{\epsilon}_{t+1}, \sigma) \equiv \mathbf{h}(\mathbf{h}(\mathbf{x}_t, \sigma) + \sigma\boldsymbol{\eta}\boldsymbol{\epsilon}_{t+1}, \sigma)$ and so forth. Pötscher and Prucha (1997) show that the system in (4) is stable if $\mathbf{h}^{(k)}$ is contracting, which is a much weaker condition than requiring $\mathbf{h}(\mathbf{x}_t, \sigma)$ to be contracting. Two sufficient conditions ensure that the contraction property holds for $\mathbf{h}^{(k)}$. First, there must exist an integer $k \geq 1$ at which

$$\sup\left\{\left|\text{stac}_{j=1}^{n_x}\left[\mathbf{i}_j' \frac{\partial \mathbf{h}^{(k)}}{\partial \mathbf{x}'}\left(\mathbf{x}^j, \{\boldsymbol{\epsilon}_d^j\}_{d=1}^{k-1}, \sigma\right)\right]\right|\right\} < 1, \tag{11}$$

given $\mathbf{x}^j \in \mathcal{X}_x$ and $\boldsymbol{\epsilon}^j \in \mathcal{X}_\epsilon$. Here, $\frac{\partial \mathbf{h}^{(k)}}{\partial \mathbf{x}'}(\mathbf{x}^j, \{\boldsymbol{\epsilon}_d^j\}_{d=1}^{k-1}, \sigma)$ is an $n_x \times n_x$ Jacobian matrix evaluated at $(\mathbf{x}^j, \{\boldsymbol{\epsilon}_d^j\}_{d=1}^{k-1})$, and $|\mathbf{A}|$ denotes the norm given by the square root of the largest eigenvalue of the matrix product $\mathbf{A}'\mathbf{A}$. The vector $\mathbf{i}_j$ is the j'th column in the $n_x \times n_x$ identity matrix, and the stac-operator creates a matrix using the rows shown as arguments to the operator.[4] Hence, the condition in (11) states that for a sufficiently large integer k,

[4]For instance, let $\mathbf{a}_j$ denote the j'th row of an $m \times n$ matrix $\mathbf{A}$, then $\text{stac}_{j=1}^{m}\mathbf{a}_j = \mathbf{A}$. The stac-operator is used in (11) to allow rows in $\partial \mathbf{h}^{(k)}/\partial \mathbf{x}'$ to be evaluated at different points, as indicated by the superindex j on the arguments at which $\partial \mathbf{h}^{(k)}/\partial \mathbf{x}'$ is evaluated.

the largest norm of $\partial \mathbf{h}^{(k)}/\partial \mathbf{x}'$ must be strictly smaller than one for *all* values of $\mathbf{x}_t$ and $\boldsymbol{\epsilon}_t$ in their feasible domains.

The second condition is much weaker than (11) and given by

$$\sup\left\{\left|\frac{\partial \mathbf{h}^{(k)}}{\partial \boldsymbol{\epsilon}_l'}\left(\mathbf{x}, \{\boldsymbol{\epsilon}_d\}_{d=1}^{k-1}, \sigma\right)\right|\right\} < \infty, \tag{12}$$

where $\mathbf{x} \in \mathcal{X}_x$ and $\boldsymbol{\epsilon}_d \in \mathcal{X}_\epsilon$ for $l = 1, 2, \ldots, k-1$. It is clear that this second condition holds for basically all smooth approximations to DSGE models if $\mathbf{x}_t$ and $\boldsymbol{\epsilon}$ are finite. We therefore focus on (11) in our subsequent discussion and leave (12) as a technical regularity condition.

3.2 *Stability of the extended perturbation approximation*

Before analyzing the stability of extended perturbation, it is useful to study the stability properties of the certainty equivalent solution. For the considered class of DSGE models in Section 2.1, the imposed assumptions ensure stability under certainty equivalence, as all state variables (in the absence of shocks) return to the steady state when time tends to infinity.[5] This stability requirement means that the state process under certainty equivalence $\mathbf{x}_t^{\mathrm{CE}}$ is stable, where $\mathbf{x}_t^{\mathrm{CE}}$ evolves as $\mathbf{x}_{t+1}^{\mathrm{CE}} = \mathbf{h}^{\mathrm{CE}}(\mathbf{x}_t^{\mathrm{CE}}) + \sigma\boldsymbol{\eta}\boldsymbol{\epsilon}_{t+1}$. In other words, $\mathbf{h}^{\mathrm{CE}}$ satisfies condition (11) and explosive sample paths for $\mathbf{x}_t^{\mathrm{CE}}$ do not appear.

We next analyze the stability of extended perturbation when gradually increasing the approximation order, that is, the Taylor series expansion of $\mathbf{h}^{\mathrm{stoch}}$. For this analysis, it is useful to write the extended perturbation approximation as $\mathbf{x}_{t+1} = \mathbf{h}^{\mathrm{ExPer}}(\mathbf{x}_t, \sigma) + \sigma\boldsymbol{\eta}\boldsymbol{\epsilon}_{t+1}$, where $\mathbf{h}^{\mathrm{ExPer}}(\mathbf{x}_t, \sigma) \equiv \mathbf{h}^{\mathrm{CE}}(\mathbf{x}_t) + \hat{\boldsymbol{h}}^{\mathrm{stoch}}(\mathbf{x}_t, \sigma)$. In a first-order approximation, there is no correction for uncertainty because $\hat{\boldsymbol{h}}^{\mathrm{stoch}} = \mathbf{0}$, meaning that extended perturbation reduces to the stable certainty equivalent solution.

In a second-order approximation, there is a constant correction for uncertainty as $\hat{\boldsymbol{h}}^{\mathrm{stoch}} = \frac{1}{2}\mathbf{h}_{\sigma\sigma}\sigma^2$. This means that partial derivatives of $\mathbf{h}^{\mathrm{ExPer}}$ with respect to the state variables are equal to those of $\mathbf{h}^{\mathrm{CE}}$ for all values of $\mathbf{x}_t$, implying that the stability condition (11) also holds for $\mathbf{h}^{\mathrm{ExPer}}$. Thus, extended perturbation at second order guarantees a stable approximation, because the uncertainty correction only recenters the stable certainty equivalent solution.

In a third-order approximation, the uncertainty correction is a linear function of the state variables as shown in (8). This implies that partial derivatives of $\mathbf{h}^{\mathrm{ExPer}}$ differ from those of $\mathbf{h}^{\mathrm{CE}}$ and the stability condition (11) cannot be guaranteed to hold for $\mathbf{h}^{\mathrm{ExPer}}$ due to the local nature of the approximation to $\mathbf{h}^{\mathrm{stoch}}$. In other words, the process for $\mathbf{x}_t$ does not necessarily inherit stability from the certainty equivalent solution, because the term $\mathbf{h}_{\sigma\sigma\mathbf{x}}$ may generate instability if the linear approximation of $\mathbf{h}^{\mathrm{stoch}}$ is insufficiently accurate.

[5] Stability under certainty equivalence is a standard assumption to impose, as it is routinely exploited at the steady state (through the Blanchard–Kahn condition) to derive linearized solutions to DSGE models and, therefore, also higher-order perturbation approximations. In any case, the stability assumption under certainty equivalence can be tested using the procedure in Boucekkine (1995).

Going beyond third order, the stochastic component of the policy function is approximated more accurately around the steady state by higher-order polynomials as shown in (9) and (10). However, as at third order, we cannot guarantee a stable approximation, because partial derivatives of $\hat{\boldsymbol{h}}^{\text{stoch}}$ may violate the stability condition in (11) for $\mathbf{h}^{\text{ExPer}}$ although satisfied for $\mathbf{h}^{\text{CE}}$.

3.3 *Testing for stability*

Given that extended perturbation does not necessarily provide a stable approximation, it seems useful to have a test to determine if a given approximation is stable. The test we propose applies to any approximation of DSGE models within the considered class and builds on two simplifying assumptions. We first propose to only evaluate (11) on a sparse grid containing extreme state values, as unstable dynamics are most likely to appear at such points given the local nature of our approximation to $\mathbf{h}^{\text{stoch}}$. To construct the grid, let $\mathcal{S}_i = \{l_i^x, u_i^x\}$ for $i = 1, 2, \ldots, n_x$ contain the lower bound l_i^x and the upper bound u_i^x of the ith state variable. The values of $\{l_i^x, u_i^x\}_{i=1}^{n_x}$ should cover the region where the approximation is used.[6] We then form the Cartesian set $\mathcal{S}_x \equiv \prod_{i=1}^{n_x} \mathcal{S}_i$, which has 2^{n_x} elements. Our second simplifying assumption is to only consider the stability condition in (11) when the rows in $\partial \mathbf{h}^{(k)}/\partial \mathbf{x}'$ are evaluated at the same point.[7]

Given these simplifying assumptions, the stability condition in (11) reduces to the testable requirement that $\mathbf{h}^{(k)}$ is contracting if there exists an integer $k \geq 1$ such that

$$\max\left\{\left|\frac{\partial \mathbf{h}^{(k)}}{\partial \mathbf{x}'}\left(\mathbf{x}, \left\{\boldsymbol{\epsilon}_d^{(v)}\right\}_{d=1}^{k-1}, \sigma\right)\right|, \text{for all } \mathbf{x} \in \mathcal{S}_x \text{ and } v = 1, 2, \ldots, \mathcal{M}\right\} < 1. \tag{13}$$

Here, each point in $\mathcal{S}_x$ is evaluated using $\mathcal{M}$ sample paths of the innovations $\{\boldsymbol{\epsilon}_d^{(v)}\}_{d=1}^{k-1}$ to avoid that a nonstable system satisfies the contraction condition just because of a "fortunate" sample path for the innovations. The test may be carried out for different values of k and $\mathcal{M}$. Some guidance on a reasonable value of k may be obtained by implementing the test on a stable linear solution. We generally recommend using a fairly large value of k, say 500 or 1000, because it is easier for $\mathbf{h}(\mathbf{x}_t, \sigma)$ to satisfy the contraction property when iterated many periods forward in time.[8]

Another method commonly used to detect unstable approximations is to simulate multiple sample paths and see if any of these simulations explode. Compared to this brute force approach, our stability test is computationally less demanding and, more importantly, allows the researcher to locate critical state configurations where the approximation might explode. As we will show below in Section 4.3, such information is

[6] Guidance on how to set these bounds may be obtained from unconditional moments or extreme values in a simulated sample using the extended perturbation approximation.

[7] At the expense of increasing the computational cost of the test, it is obvious that a finer grid for the state variables may be considered and that rows in $\partial \mathbf{h}^{(k)}/\partial \mathbf{x}'$ could be evaluated at different points.

[8] Given that the Jabocian $\partial \mathbf{h}^{(k)}/\partial \mathbf{x}'$ is computed by numerical differentiation, the most efficient implementation of the test is to evaluate $|\partial \mathbf{h}^{(j_{\mathbf{x}})}/\partial \mathbf{x}'|$ by gradually increasing $j_{\mathbf{x}}$ and then stop for a given $\mathbf{x} \in \mathcal{S}_x$ when the condition is met, even though $j_{\mathbf{x}}$ may be less than a predetermined value of k. If $\max\{\{j_{\mathbf{x}}\}_{\mathbf{x} \in \mathcal{S}_x}\} < k$, then the stability condition in (13) is satisfied.

valuable because it may allow the researcher to understand *why* a given approximation is unstable.

3.4 *Inducing stability*

Section 3.2 showed that computing $\hat{\boldsymbol{h}}^{\text{stoch}}(\mathbf{x}_t, \sigma)$ by standard perturbation may lead to instability of extended perturbation beyond second order. If a given approximation is unstable, the most natural solution is to increase the approximation order for $\hat{\boldsymbol{h}}^{\text{stoch}}(\mathbf{x}_t, \sigma)$ until the condition in (13) is satisfied. If this is not possible due to computational constraints, another way to ensure stability is to downscale $\hat{\boldsymbol{h}}^{\text{stoch}}(\mathbf{x}_t, \sigma)$ as we move away from steady state, where this local uncertainty correction may be less accurate. Inspired by the work of Blasques and Nientker (2020), we consider the function $\Phi(\mathbf{x}_t) \equiv e^{-\tau \times D(\mathbf{x}_t)}$, where $\tau \geq 0$ and $D(\mathbf{x}_t) = \sum_{i=1}^{n_x}(x_{i,t} - x_{i,\text{ss}})^2$ measures the distance of $\mathbf{x}_t$ from $\mathbf{x}_{\text{ss}}$. Thus, if extended perturbation is unstable, then we suggest to use the modified state transition

$$\mathbf{x}_{t+1} = \mathbf{h}^{\text{CE}}(\mathbf{x}_t) + \hat{\boldsymbol{h}}^{\text{stoch}}(\sigma) + \big(\hat{\boldsymbol{h}}^{\text{stoch}}(\mathbf{x}_t, \sigma) - \hat{\boldsymbol{h}}^{\text{stoch}}(\sigma)\big)\Phi(\mathbf{x}_t) + \sigma\boldsymbol{\eta}\boldsymbol{\epsilon}_{t+1}, \tag{14}$$

where variation in $\hat{\boldsymbol{h}}^{\text{stoch}}(\mathbf{x}_t, \sigma)$ from the constant risk correction $\hat{\boldsymbol{h}}^{\text{stoch}}(\sigma)$ is downscaled to ensure stability. As an example of (14), consider a third-order approximation where $\hat{\boldsymbol{h}}^{\text{stoch}}(\sigma) = \frac{1}{2}\mathbf{h}_{\sigma\sigma}\sigma^2 + \frac{1}{6}\mathbf{h}_{\sigma\sigma\sigma}\sigma^3$, implying that only $\hat{\boldsymbol{h}}^{\text{stoch}}(\mathbf{x}_t, \sigma) - \hat{\boldsymbol{h}}^{\text{stoch}}(\sigma) = \frac{3}{6}\mathbf{h}_{\sigma\sigma\mathbf{x}}\mathbf{x}_t$ is scaled by $\Phi(\mathbf{x}_t)$. It is easy to see that $\Phi(\mathbf{x}_t)$ in (14) does not affect the level or the curvature of the state transition function $\mathbf{h}(\mathbf{x}_t)$ at the steady state, implying that the modified state transition in (14) also matches the curvature of the true solution at the steady state. In comparison to the work of Blasques and Nientker (2020), we only downscale $\hat{\boldsymbol{h}}^{\text{stoch}}(\mathbf{x}_t, \sigma) - \hat{\boldsymbol{h}}^{\text{stoch}}(\sigma)$, whereas they also downscale all nonlinear terms in the Taylor series expansions of $\mathbf{h}^{\text{CE}}(\mathbf{x}_t)$ to ensure stability of standard perturbation. As for the value of τ, one possibility is to select the smallest τ that satisfies the stability test in Section 3.3. A somewhat more computational demanding alternative is to follow the procedure in Blasques and Nientker (2020) and set τ to minimize the model's Euler equation errors on some predefined grid.

4. A New Keynesian model

We next study the accuracy and stability of standard and extended perturbation using a New Keynesian model with price stickiness as in Calvo (1983). Two reasons motivate our choice of model. First, the New Keynesian model with Calvo pricing is one of the most popular DSGE models in the literature. Second, and perhaps somewhat surprisingly, some dimensions of this New Keynesian model are highly nonlinear even for a standard calibration. The strong nonlinearities in the model also imply that standard perturbation easily generate explosive dynamics, showing that unstable approximations are not only seen at extreme calibrations, as found in Haan and De Wind (2012) for the neoclassical growth model. We proceed by describing the New Keynesian model in Section 4.1, before studying accuracy in Section 4.2 and stability in Section 4.3 for approximations up to fifth order.

4.1 *Model description*

A representative household maximizes

$$\mathcal{U}_t = E_t \sum_{l=0}^{\infty} \beta^l \left(\frac{c_{t+l}^{1-\phi_2}}{1-\phi_2} + \phi_0 \frac{(1-h_{t+l})^{1-\phi_1}}{1-\phi_1} \right), \tag{15}$$

where c_t is consumption and h_t is labor supply. In addition to a no-Ponzi-game condition, the optimization is subject to the real budget constraint

$$c_t + b_t + i_t = h_t w_t + r_t^k k_t + \frac{R_{t-1} b_{t-1}}{\pi_t} + \text{div}_t,$$

where resources are allocated to consumption, one-period nominal bonds b_t, and investment i_t. Letting w_t denote the real wage and r_t^k the rental rate of capital k_t, the household receives (i) labor income $w_t h_t$, (ii) income from capital services sold to firms $r_t^k k_t$, (iii) payoffs from bonds purchased in the previous period $R_{t-1} b_{t-1}/\pi_t$, and (iv) dividends from firms div_t. Here, $\pi_t \equiv P_t/P_{t-1}$ is gross inflation and R_t is the gross nominal interest rate. The optimization of (15) is also subject to the law of motion for capital $k_{t+1} = (1-\delta)k_t + i_t - \frac{\kappa}{2}(\frac{i_t}{k_t} - \psi)^2 k_t$, where $\kappa \geq 0$ introduces capital adjustment costs based on i_i/k_t as in Jermann (1998). The constant ψ ensures that these adjustment costs are zero in the steady state.

We consider a perfectly competitive representative firm that produces final output y_t using the intermediate goods $y_{i,t}$ and the production function $y_t = (\int_0^1 y_{i,t}^{(\eta-1)/\eta}\, di)^{\eta/(\eta-1)}$ with $\eta > 1$. This generates the demand function $y_{i,t} = (\frac{P_{i,t}}{P_t})^{-\eta} y_t$ for the ith input $y_{i,t}$, where $P_t = [\int_0^1 P_{i,t}^{1-\eta}\, di]^{1/(1-\eta)}$ denotes the aggregate price level and $P_{i,t}$ is the price of the ith good.

The intermediate goods are produced by monopolistic competitors using the production function $y_{i,t} = a_t k_{i,t}^{\theta} h_{i,t}^{1-\theta}$, where technology a_t evolves as $\log a_{t+1} = \rho_a \log a_t + \sigma_a \epsilon_{a,t+1}$ with $\epsilon_{a,t+1} \sim \mathcal{NID}(0,1)$. The ith firm sets $P_{i,t}$, $h_{i,t}$, and $k_{i,t}$ by maximizing the present value of dividends. Beyond a no-Ponzi-game condition, the firm must satisfy demand for the ith good. When setting prices, we follow Calvo (1983) and assume that only a fraction $\alpha \in [0,1)$ of firms set their prices optimally, with the remaining firms letting $P_{i,t} = P_{i,t-1}$.

Finally, monetary policy is specified by a standard Taylor rule

$$\log(R_t/R_{\text{ss}}) = \rho \log(R_{t-1}/R_{\text{ss}}) + (1-\rho)(\kappa_\pi \hat{\pi}_t + \kappa_y \hat{y}_t), \tag{16}$$

where the policy rate is determined based on a desire to close the inflation gap $\hat{\pi}_t \equiv \log(\pi_t/\pi_{\text{ss}})$ and the output gap $\hat{y}_t \equiv \log(y_t/y_{\text{ss}})$, subject to smoothing changes in the policy rate with $\rho \in [0,1)$.

We adopt a relative standard parametrization for a quarterly model in Table 2. Households have an intertemporal elasticity of substitution of 0.5 ($\phi_2 = 2$), allocate one-third of their time endowment to labor in steady state ($h_{\text{ss}} = 0.33$), and have a Frisch labor supply elasticity of one ($\phi_1 = 2$). Firms reset prices once a year on average ($\alpha = 0.75$) and impose an average markup of 20% ($\eta = 6$). The main objective of the central bank

TABLE 2. The structural parameters.

$\beta = 0.99$	$\alpha = 0.75$
$h_{\text{ss}} = 0.33$	$\rho = 0.80$
$\phi_1 = 2.00$	$\kappa_\pi = 1.50$
$\phi_2 = 2.00$	$\kappa_y = 0.125$
$\kappa = 2.00$	$\rho_a = 0.98$
$\delta = 0.025$	$\sigma_a = 0.006$
$\theta = 0.36$	$\pi_{\text{ss}} = \begin{cases} 1.000 & \text{case } \mathcal{A} \\ 1.005 & \text{case } \mathcal{B} \end{cases}$
$\eta = 6.00$	

is to stabilize inflation ($\kappa_\pi = 1.5$, $\kappa_y = 0.125$), subject to smoothing changes in the policy rate ($\rho = 0.8$). We consider two values for the steady-state inflation rate π_{ss}. The first specification (case $\mathcal{A}$) imposes $\pi_{\text{ss}} = 1.00$, while the second specification (case $\mathcal{B}$) introduces trend inflation by letting $\pi_{\text{ss}} = 1.005$ to get an annual steady-state inflation rate of 2%.

4.2 *Accuracy analysis*

Accuracy of standard and extended perturbation are evaluated by computing unit-free Euler errors $\{E_{i,t}\}_{i=1}^n$ at state $\mathbf{x}_t$ for all $n = 16$ equations in the New Keynesian model.[9] We report the mean absolute error MAE$= \log_{10}(\sum_{t=1}^T \sum_{i=1}^n |E_{i,t}|/(nT))$, the root mean squared error RMSE$= \log_{10}(\sqrt{\sum_{t=1}^T \sum_{i=1}^n E_{i,t}^2/(nT)})$, and the maximum error MaxE$=$ $\log_{10}(\max(\{\{E_{i,t}\}_{i=1}^n\}_{t=1}^T))$, which are transformed by the 10-base logarithm.

We first consider specification $\mathcal{A}$ without trend inflation in Table 3. Panel I evaluates accuracy on a grid constructed using five points uniformly spaced along each dimension of the four-dimensional state space, giving a total of $5^4 = 625$ points. The bounds for the first three state variables (i.e., R_{t-1}, k_t, and a_t) in this grid range from -2 to $+2$ standard deviations in a log-linearized solution. For the fourth state variable, which is the log-transformed price dispersion s_t linked to the Calvo pricing, the considered range is from 0.00 to 0.04. We find that standard perturbation performs well on this grid with progressively smaller errors as measured by the MAE and RMSE when increasing the approximation order. For instance, the MAE is -2.24 (or $10^{-2.24} = 0.0058$) at first order and -3.20 (or $10^{-3.20} = 0.0006$) at fifth order. However, extended perturbation is even more accurate and outperforms standard perturbation at every approximation order as indicated by the bold figures in Table 3. This satisfying performance of extended perturbation is due to the high accuracy of the certainty equivalent solution by the extended path (i.e., extended perturbation at first order), which even outperforms standard perturbation at fifth order. We also note that the MAE becomes smaller for extended perturbation with a higher approximation order, showing that increasing the approximation order of $\mathbf{g}_t^{\text{stoch}}$ and $\mathbf{h}_t^{\text{stoch}}$ help to improve accuracy.

[9]The Appendix in the Online Supplementary Material contains a summary of these 16 equations expressed in unit-free terms.

TABLE 3. Accuracy of the New Keynesian model. This table reports the accuracy of standard and extended perturbation for the two considered calibrations in Table 2 when the states are obtained from a grid (panels I and III) or a simulated sample (panels II and IV). The grid is constructed using 5 points uniformly spaced along each dimension of the state space, giving $5^4 = 625$ grid points. The bounds for the interest rate, capital, and technology in the grid range from -2 to +2 standard deviations in a log-linearized solution. For the price dispersion index, the grid ranges from 0.00 to 0.04 in case $\mathcal{A}$ and from 0.00 to 0.10 in case $\mathcal{B}$. The states in the simulated sample are obtained from simulating extended perturbation at fifth order for 5000 observations following a burn-in of 100 observations. The conditional expectations in the Euler equations are evaluated using Gauss–Hermite quadratures with five points. Bold figures highlight the approximation with the best performance along a given accuracy measure.

		Standard Perturbation			Extended Perturbation		
		MAE	RMSE	MaxE	MAE	RMSE	MaxE
Case $\mathcal{A}$	No trend inflation ($\pi_{ss} = 1.00$)						
	Panel I: On grid						
	1st order	−2.24	−1.72	−0.55	**−3.39**	**−2.63**	**−1.25**
	2nd order	−2.59	−2.03	−0.61	**−3.42**	**−2.63**	**−1.23**
	3rd order	−2.82	−2.13	−0.52	**−3.44**	**−2.63**	**−1.22**
	4th order	−3.04	−2.22	−0.52	**−3.45**	**−2.63**	**−1.21**
	5th order	−3.20	−2.30	−0.53	**−3.46**	**−2.63**	**−1.19**
	Panel II: On simulated sample						
	1st order	−2.84	−2.20	**−0.77**	**−3.19**	**−2.23**	−0.71
	2nd order	−2.87	−1.98	−0.64	**−3.21**	**−2.23**	**−0.71**
	3rd order	−2.59	−1.55	0.00	**−3.23**	**−2.23**	**−0.56**
	4th order	−0.93	1.37	3.80	**−3.24**	**−2.24**	**−0.56**
	5th order	0.07	2.22	4.62	**−3.25**	**−2.24**	**−0.56**
Case $\mathcal{B}$	With trend inflation ($\pi_{ss} = 1.005$)						
	Panel III: On grid						
	1st order	−2.10	−1.59	−0.45	**−2.99**	**−2.29**	**−0.97**
	2nd order	−2.19	−1.49	0.22	**−2.99**	**−2.29**	**−0.97**
	3rd order	−1.69	0.11	2.11	**−2.99**	**−2.29**	**−0.96**
	4th order	−0.76	1.06	2.99	**−2.98**	**−2.29**	**−0.95**
	5th order	−1.82	−0.59	1.18	**−2.98**	**−2.29**	**−0.94**
	Panel IV: On simulated sample						
	1st order	−2.67	**−2.11**	**−0.78**	**−2.90**	−2.09	−0.59
	2nd order	−2.58	−1.87	**−0.63**	**−2.92**	**−2.10**	−0.58
	3rd order	−2.26	−1.41	0.16	**−2.93**	**−2.10**	**−0.58**
	4th order	−0.22	2.21	4.66	**−2.93**	**−2.11**	**−0.58**
	5th order	5.03	7.47	9.92	**−2.94**	**−2.12**	**−0.58**

Panel II evaluates accuracy on a simulated sample of 5000 states from extended perturbation at fifth order, which we then use to study the accuracy of both perturbation methods at the various approximation orders. We first note that standard perturbation at first order is very accurate with a MAE of −2.84 and a RMSE of −2.20, but that performance deteriorates fairly rapidly when increasing the approximation order. For instance, the MAE is −2.59 at third order and only 0.07 at fifth order. Unreported results

reveal that a higher approximation order gives smaller errors close to the steady state, but also much larger errors far from the steady state. This is particularly evident when the price dispersion exceeds the range [0.00, 0.04] used for the grid in Panel I, although it only happens in 6% of our simulated sample. On the other hand, extended perturbation remains very accurate and clearly outperforms standard perturbation with a MAE of -3.19 at first order and a MAE of -3.25 at fifth order. This improved performance must be due to the more accurate certainty equivalence solution by the extended path, which implicitly adds the higher order terms $\sum_{k=m+1}^{\infty} \frac{g(x_{\text{ss}},0)_{\mathbf{x}^k}}{k!}(x_t - x_{\text{ss}})^k$ to the standard m'th order Taylor-series approximation as shown above. Given that there is no singularity in the New Keynesian model, it also means that standard perturbation would eventually match the performance of extended perturbation for a sufficiently high approximation order and eventually converge to the true solution on a very large domain.

For specification $\mathcal{B}$ with trend inflation, standard perturbation at first order does fairly well, but the accuracy deteriorates when increasing the approximation order, both on the grid (Panel III) and in the simulated sample (Panel IV). As for specification $\mathcal{A}$ without trend inflation, unreported results show that a higher approximation order gives smaller errors close to the steady state, but also much larger errors when the price dispersion is far from steady state. This happens more often with trend inflation, where the price dispersion exceeds 0.04 in 15% of our simulated sample, whereas the corresponding figure is 6% for specification $\mathcal{A}$. In contrast, extended perturbation remains highly accurate with trend inflation and clearly does better than standard perturbation for nearly all orders of the approximation.

4.3 *Stability analysis*

This section uses our stability test from Section 3.3 to study the dynamic properties of standard and extended perturbation and to understand *why* a given approximation may be unstable. To run the stability test for the New Keynesian model with specification $\mathcal{A}$, we first construct the set $\mathcal{S}_x$ with extreme state configurations. Using the conventional notation, where a "hat" denotes percentage deviation from steady state, the bounds for $\hat{a}_t$ are given by ± 3 standard deviations of technology. For the two endogenous states $(\hat{R}_{t-1}, \hat{k}_t)$, the bounds are set to ± 4 standard deviations in a log-linearized solution, and hence slightly wider than for technology to account for effects of nonlinearities. The price dispersion s_t is zero in steady state and constant in a log-linearized solution without trend inflation, and we therefore use a simulated sample path of extended perturbation to guide our bounds of 0.00 and 0.08 for $\hat{s}_t$. Using this specification of $\mathcal{S}_x$, we find that standard perturbation at second order passes the stability test with $\mathcal{M} = 50$, but not at third, fourth, and fifth order. This is because the state dynamics at several points diverge when iterated forward in time to compute the Jacobian $\partial \mathbf{h}^{(k)}/\partial \mathbf{x}'$, which at these points display explosive eigenvalues. In contrast, extended perturbation at third, fourth, and fifth order pass the stability test with $\mathcal{M} = 50$. This reveals that the explosive behavior of standard perturbation beyond second order must be due to approximation errors in the certainty equivalent component of the policy function, as the two perturbation methods rely on the same approximation to the stochastic part of the policy function.

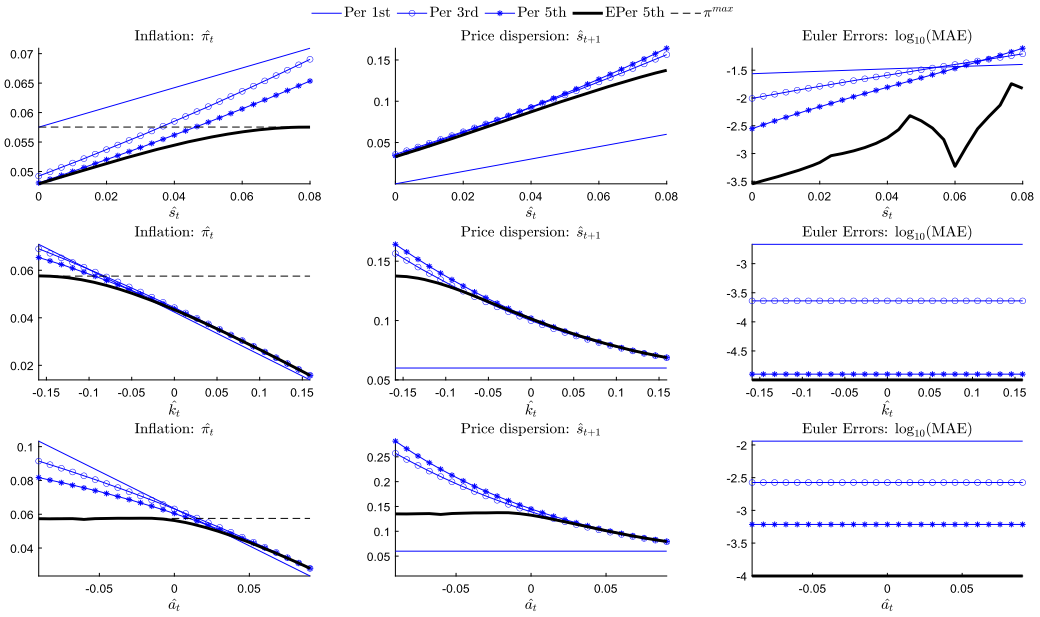

FIGURE 1. Inflation and price dispersion at extreme state configurations. The top row shows inflation $\hat{\pi}_t$, the price dispersion next period $\hat{s}_{t+1}$, and the mean absolute Euler errors across the 16 equations in the model when increasing $\hat{s}_t$. The middle and the bottom rows show the corresponding plots when changing the capital stock $\hat{k}_t$ and the technology $\hat{a}_t$ level, respectively. Unless stated otherwise, the state variables that are not plotted are assigned the following values: (i) the capital stock $\hat{k}_t$ and the lagged interest rate $\hat{R}_{t-1}$ equals -4 standard deviations in a log-linear solution, (ii) the technology level $\hat{a}_t$ equals -3 standard deviations, and (iii) the price dispersion $\hat{s}_t$ is 0.08. All structural parameters attain the values provided in Table 2 for specification $\mathcal{A}$. In the legends, "Per" denotes standard perturbation and "EPer" is an abbreviation for extended perturbation.

For specification $\mathcal{B}$ with trend inflation, the same procedure as in specification $\mathcal{A}$ is used to set the bounds for $\hat{R}_{t-1}$, $\hat{k}_t$, and $\hat{a}_t$ in the stability test, while the upper bound of $\hat{s}_t$ is increased to 0.2 to better capture the more frequent extreme values of $\hat{s}_t$ with trend inflation. We find that standard perturbation is unstable at second order and beyond, whereas extended perturbation at third, fourth, and fifth order once again passes the stability test with $\mathcal{M} = 50$.

An inspection of the $2^4 = 16$ extreme state configurations in $\mathcal{S}_x$ reveals that a high price dispersion, a low capital stock, and a low technology level may lead to explosive dynamics for standard perturbation. To understand why, consider specification $\mathcal{A}$ in Figure 1, where we plot $\hat{\pi}_t$ and $\hat{s}_{t+1}$ as we change some of the state variables, while the remaining states attain their extreme values from $\mathcal{S}_x$, where the approximation is unstable. The top left chart shows that inflation $\hat{\pi}_t$ in standard perturbation at third and fifth order (the lines marked with circles and stars) increase sharply for higher values of $\hat{s}_t$. This increase in inflation leads to even higher values of the price dispersion in the next period $\hat{s}_{t+1}$, as seen in the middle chart at the top. A higher value of $\hat{s}_{t+1}$ then

increases $\hat{\pi}_{t+1}$, which in turn increases $\hat{s}_{t+2}$ and so on. That is, standard perturbation generates a price-inflation spiral, which eventually may lead to explosive dynamics. The inflation dynamics from extended perturbation (the thick black line) is radically different, because inflation converges to an upper threshold (the dotted line) for large values of $\hat{s}_t$. This implies that $\hat{s}_{t+1}$ increases more slowly with higher values of $\hat{s}_t$, which helps to improve accuracy, as shown by the top right chart in Figure 1. To see where this threshold comes from, recall that the aggregate price level $P_t = [\int_0^1 P_{i,t}^{1-\eta}\, di]^{1/(1-\eta)}$ reduces to $P_t^{1-\eta} = (1-\alpha)\tilde{P}_t^{1-\eta} + \alpha P_{t-1}^{1-\eta}$ with Calvo pricing, or equivalently

$$1 = (1-\alpha)\left(\frac{\tilde{P}_t}{P_t}\right)^{1-\eta} + \alpha\left(\frac{1}{\pi_t}\right)^{1-\eta}, \tag{17}$$

where $\tilde{P}_t$ denotes the new price of firms resetting their prices in period t. Consider the case where $\tilde{P}_t$ is very large, for instance, due to a high price dispersion, a low capital stock, or a low technology level. The constant elasticity of substitution (CES) function then implies that $(1-\alpha)(\tilde{P}_t/P_t)^{1-\eta}$ becomes very small as $\eta > 1$. In the limit when $\tilde{P}_t$ becomes arbitrarily large, this term converges to zero and (17) implies the upper bound $\pi^{\max} = \alpha^{1/(1-\eta)}$ for inflation. Intuitively, consumers simply substitute away from the very expensive goods to the cheaper goods whose prices have not yet been updated, and this ensures that inflation remains bounded.[10] This upper bound on inflation appears to have been largely ignored in the literature (perhaps due to the widespread use of a log-linearization), but our results show that it is important to ensure stability of a nonlinear approximation to the New Keynesian model.

The charts in the middle and the bottom row of Figure 1 plot $\hat{\pi}_t$, $\hat{s}_{t+1}$, and Euler errors when we instead vary the capital stock and the technology level, respectively. Standard perturbation is again unable to account for the upper bound on inflation, and this leads to a high value of the price dispersion in the next period $\hat{s}_{t+1}$, which then increases $\hat{\pi}_{t+1}$, and hence starts the price-inflation spiral. Extended perturbation does not suffer from this shortcoming, as it perfectly accounts for the upper bound on inflation, and hence provides a more accurate approximation, as seen from the Euler errors in the charts to the right in Figure 1.

Thus, our analysis reveals that a price-inflation spiral plays an important role in generating explosive dynamics in the New Keynesian model when using standard perturbation beyond first order. This means that various modifications to the model that imply less variation in inflation and the price dispersion should reduce the probability of starting the price-inflation spiral and generating explosive sample paths. Omitting trend inflation is one way to reduce the variability in s_t as shown above. Another possibility is to follow Christiano, Eichenbaum, and Evans (2005) and adopt price-indexation for nonoptimizing firms by letting $P_{i,t} = P_{i,t-1}\pi_{t-1}$, as this also reduces the variability of s_t.

[10]We are grateful to an anonymous referee for helping us realize the economic intuition behind this result. If we accommodate price indexation of the nonoptimizing firms, then (17) becomes $1 = (1-\alpha)(\tilde{P}_t/P_t)^{1-\eta} + \alpha(\pi_{t-1}/\pi_t)^{1-\eta}$, which implies an upper bound on the growth rate of inflation, that is, $\pi_t/\pi_{t-1} = \alpha^{1/(1-\eta)}$.

One may also consider a central bank that has a strong preference for closing the inflation gap (i.e., a high κ_π), as this reduces the variation in inflation and hence makes it less likely that inflation hits its upper bound.

5. Computational issues related to extended perturbation

This section deals with the computational issues related to extended perturbation. Section 5.1 presents an efficient procedure to compute the certainty equivalent solution, and Section 5.2 discusses issues related to including extended perturbation in an estimation routine.

5.1 *The certainty equivalent component of the policy function*

The first step of extended perturbation involves computing derivatives of the model at the steady state. The computational requirement of this step depends on the size of the model and the considered order of approximation, as shown in Section 2.4. The second step of extended perturbation is to obtain the certainty equivalent solution. Given the state $\mathbf{x}_t$, the extended path computes $\mathbf{g}^{\text{CE}}(\mathbf{x}_t)$ and $\mathbf{h}^{\text{CE}}(\mathbf{x}_t)$ by truncating the infinite system in (1) at period $t+N$, where a terminal value for the controls $\mathbf{y}_{t+N}$ is imposed. In the absence of uncertainty, this implies the finite dimensional system

$$\begin{aligned}
\mathbf{f}(\mathbf{x}_t, \mathbf{x}_{t+1}, \mathbf{y}_t, \mathbf{y}_{t+1}) &= \mathbf{0}_{n\times 1},\\
\mathbf{f}(\mathbf{x}_{t+1}, \mathbf{x}_{t+2}, \mathbf{y}_{t+1}, \mathbf{y}_{t+2}) &= \mathbf{0}_{n\times 1},\\
&\dots\\
\mathbf{f}(\mathbf{x}_{t+N-1}, \mathbf{x}_{t+N}, \mathbf{y}_{t+N-1}, \mathbf{y}_{t+N}) &= \mathbf{0}_{n\times 1},
\end{aligned} \tag{18}$$

with a total of nN equilibrium conditions. The initial state $\mathbf{x}_t$ and $\mathbf{y}_{t+N}$ are known by assumption, and this allows the extended path to find the nN unknowns in $\{\mathbf{x}^{\text{EP}}_{t+i}\}_{i=1}^{N}$ and $\{\mathbf{y}^{\text{EP}}_{t+i}\}_{i=0}^{N-1}$ by solving the fixed-point problem in (18), with $\mathbf{g}^{\text{CE}}(\mathbf{x}_t) = \mathbf{y}^{\text{EP}}_t$ and $\mathbf{h}^{\text{CE}}(\mathbf{x}_t) = \mathbf{x}^{\text{EP}}_{t+1}$. This fixed-point problem is typically solved within a few iterations using the Newton–Raphson algorithm of Boucekkine (1995), that exploits the sparsity and recursive nature of the system in (18). When this algorithm fails to converge, we simply minimize the squared residuals for the system in (18) using a Levenberg–Marquardt-type nonlinear minimizer. Although each of these solution routines are very fast, the computational burden may nevertheless accumulate if the extended path is called repeatedly, for instance, when simulating a long sample. We therefore present several refinements of the extended path to make it more efficient.

First, the exogenous states $\{\mathbf{x}_{2,t+j}\}_{j=1}^{N}$ can be eliminated from (18), as they can be solved directly by iterating on (2). It is also desirable to eliminate any lagged control variables that appear in $\mathbf{x}_t$ from (18). For instance, in our New Keynesian model, we do not need to separately solve for the short rate $\{R_t\}_{t=1}^{N}$ in the controls and the lagged short rate $\{R_{t-1}\}_{t=1}^{N}$ in the states. By eliminating these variables, we reduce the number of equations and unknowns in (18), which make the system easier to solve. These simplifications also ensure that the extended path performs well on large models, as the

TABLE 4. Computational costs of extended perturbation. This table reports the accuracy of extended perturbation at fifth order and various execution statistics for the same grid as considered in Table 3 when using calibration $\mathcal{A}$. We let $D_{\text{ss}} = 0.00$ and $\mathcal{R} = 0$ in panel $\mathcal{A}$, while $\mathcal{R} = 0$ in panel $\mathcal{B}$. The computations are carried out on a standard desktop with Windows 10 using an Intel(R) Core(TM) i7-7600U CPU @ 280 GHz processor.

	Accuracy: MAE	Seconds	Average Iterations	Average N_t^*	Fraction by Extended Path
Panel $\mathcal{A}$					
Start = 1st	−3.45	964	3.3	200	100%
Start = 4th	−3.45	881	2.4	200	100%
Panel $\mathcal{B}$					
Start = 4th, $D_{\text{ss}} = 0.01$	−3.45	654	2.5	120	100%
Start = 4th, $D_{\text{ss}} = 0.03$	−3.45	333	2.5	72	100%
Start = 4th, $D_{\text{ss}} = 0.05$	−3.47	134	2.5	57	100%
Panel $\mathcal{C}$					
Start = 4th, $D_{\text{ss}} = 0.05$, $\mathcal{R} = 10^{-4}$	−3.46	105	1.7	64	44%
Start = 4th, $D_{\text{ss}} = 0.05$, $\mathcal{R} = 10^{-3}$	−3.45	94	2.2	71	19%

number of shocks and lagged control variables only have a small effect on its execution time.

Second, it is important to have an accurate terminal condition $\mathbf{y}_{t+N}$ and good starting values for $\{\mathbf{x}_{t+i}^{\text{EP}}, \mathbf{y}_{t+i-1}^{\text{EP}}\}_{i=1}^{N}$ to obtain fast convergence of (18). For the benchmark version of the extended path, these values are computed using a first-order approximation as in Adjemian and Juillard (2010). For extended perturbation, higher-order derivatives of the $\mathbf{g}$- and $\mathbf{h}$-functions are available, as they are required to compute the uncertainty corrections. It therefore seems natural to use these higher-order derivatives when computing $\{\mathbf{x}_{t+i}^{\text{EP}}, \mathbf{y}_{t+i-1}^{\text{EP}}\}_{i=1}^{N}$ and $\mathbf{y}_{t+N}$; however, a higher approximation order may not always improve accuracy far from the steady state if the model is very nonlinear, as shown in Section 4. The considered procedure is therefore to use a fourth-order approximation to obtain $\{\mathbf{x}_{t+i}^{\text{EP}}, \mathbf{y}_{t+i-1}^{\text{EP}}\}_{i=1}^{N}$ and $\mathbf{y}_{t+N}$, except when any of these elements are far from their steady state—say, differ more than 50% or 100% from the steady state. In these events, we simply use a first-order approximation. To illustrate the computational gain from the different starting values, let us return to the setting in Section 4.2 where the performance of extended perturbation at fifth order is evaluated on a grid. Panel $\mathcal{A}$ in Table 4 shows that it takes 964 seconds to compute the required Euler errors using starting values from a first-order approximation, where the extended path on average uses 3.3 iterations per observation. When using starting values from a fourth-order approximation, the execution time falls to 881 seconds, because the extended path on average only uses 2.4 iterations per observation. Importantly, this reduction in execution time is achieved without affecting the accuracy of the approximation.

Third, it is desirable to use a relatively short horizon N, as the system in (18) grows linearly in N. The extended path is commonly applied with the same N for all states. But, if $\mathbf{x}_t$ is close to $\mathbf{x}_{\text{ss}}$, a smaller N should be needed than when $\mathbf{x}_t$ is far from $\mathbf{x}_{\text{ss}}$. We

therefore introduce a state-dependent horizon N_t^* using the rule

$$N_t^* = \min\{N \in \mathbb{N} : \max\{|\mathbf{y}_{t+N} - \mathbf{y}_{ss}|\} \leq D_{ss} \text{ st. } N \in [N_{min}, N_{max}]\}. \tag{19}$$

That is, we select the shortest horizon where the largest element in $\mathbf{y}_{t+N}$ is within the distance D_{ss} from its steady state, subject to $N \in [N_{min}, N_{max}]$. This implies that N_t^* depends on D_{ss} and the current state $\mathbf{x}_t$ through $\mathbf{y}_{t+N}$. Standard perturbation is typically fairly accurate close to the steady state, and we therefore suggest to let $D_{ss} = 0.03$ or $D_{ss} = 0.05$ such that $\mathbf{y}_{t+N}$ is within 3% or 5%, respectively, of its steady-state value. As for the bounds, our experience is that $N_{min} = 50$ and $N_{max} = 200$ perform well and imply enough periods for $\mathbf{y}_{t+N}$ to get sufficiently close to $\mathbf{y}_{ss}$. Table 4 documents a substantial reduction in the average horizon N from 200 in Panel $\mathcal{A}$ (with $D_{ss} = 0$) to only 57 in Panel $\mathcal{B}$ with $D_{ss} = 0.05$. As a result, the execution time falls from 881 seconds in Panel $\mathcal{A}$ to just 134 seconds in Panel $\mathcal{B}$ with $D_{ss} = 0.05$, while the accuracy is basically unaffected.

Fourth, a standard perturbation approximation to $\mathbf{g}_t^{CE}$ and $\mathbf{h}_t^{CE}$ may for some states be sufficiently accurate, for instance when $\mathbf{x}_t$ is close to $\mathbf{x}_{ss}$ or when the model does not display strong nonlinearities. This suggests that we do not always have to run the extended path. To formalize this observation, let $\mathbf{y}_t^{Per} \equiv \mathbf{g}^{Per}(\mathbf{x}_t)$, $\mathbf{x}_{t+1}^{Per} \equiv \mathbf{h}^{Per}(\mathbf{x}_t)$, and $\mathbf{y}_{t+1}^{Per} = \mathbf{g}^{Per}(\mathbf{x}_{t+1}^{Per})$ denote the standard perturbation solution to the certainty equivalent part of the policy function. In the absence of uncertainty, the unit-free Euler errors for this approximation are $\mathbf{r}_t \equiv \mathbf{f}(\mathbf{x}_t, \mathbf{x}_{t+1}^{Per}, \mathbf{y}_t^{Per}, \mathbf{y}_{t+1}^{Per})$. The suggested approximation is then given by

$$\mathbf{y}_t = 1_{\{\max|\mathbf{r}_t| \leq \mathcal{R}\}} \mathbf{g}^{Per}(\mathbf{x}_t) + 1_{\{\max|\mathbf{r}_t| > \mathcal{R}\}} \mathbf{g}^{CE}(\mathbf{x}_t) + \hat{\boldsymbol{g}}^{stoch}(\mathbf{x}_t), \tag{20}$$

$$\mathbf{x}_{t+1} = 1_{\{\max|\mathbf{r}_t| \leq \mathcal{R}\}} \mathbf{h}^{Per}(\mathbf{x}_t) + 1_{\{\max|\mathbf{r}_t| > \mathcal{R}\}} \mathbf{h}^{CE}(\mathbf{x}_t) + \hat{\boldsymbol{h}}^{stoch}(\mathbf{x}_t) + \sigma \boldsymbol{\eta} \boldsymbol{\epsilon}_{t+1}, \tag{21}$$

where $\mathcal{R}$ denotes the tolerated errors in the certainty equivalent component of the policy function. That is, we use the standard perturbation solution when $\max|\mathbf{r}_t| \leq \mathcal{R}$, and only run the extended path when the perturbation solution is considered to be insufficiently accurate with $\max|\mathbf{r}_t| > \mathcal{R}$. The recommendation is to consider a value of $\mathcal{R}$ that is close to the accuracy level of extended perturbation for a given model. Hence, in our case, we consider values of $\mathcal{R} = 10^{-4}$ or $\mathcal{R} = 10^{-3}$, given that the Euler errors when accounting for uncertainty in the New Keynesian model typically are in this range (see Section 4.2). Panel $\mathcal{C}$ in Table 4 shows that this refinement further reduces the execution time to 105 seconds with $\mathcal{R} = 10^{-4}$ and to 94 seconds with $\mathcal{R} = 10^{-3}$, as we now only run the extended path for 44% and 19% of the points on the grid, respectively. Again, these improvements are obtained with no loss in accuracy. Despite these substantial reduction in execution time (from 964 to 94 seconds), we also note that extended perturbation remains computational more demanding than standard perturbation, which at fifth order only requires 0.5 seconds to evaluate these Euler errors.

To illustrate the benefit of these refinements, we consider the execution time of using extended perturbation at fifth order to simulate 2000 observations for the New Keynesian model with specification $\mathcal{A}$. It takes 227 seconds when only eliminating exogenous

states and lagged control variables with $N = 200$. Introducing the state-dependent horizon N_t^* in (19) with $D_{ss} = 0.05$ reduces the execution time to 61 seconds. If we in addition use (20) and (21) to occasionally compute the certainty equivalent component of the policy function by standard perturbation, then we reduce the execution time to 18 seconds with $\mathcal{R} = 10^{-4}$ and 12 seconds with $\mathcal{R} = 10^{-3}$.[11] In contrast, simulating 2000 observation using standard perturbation at fifth order takes 0.4 seconds.

5.2 *Estimation*

The previous section showed that the computational costs of extended perturbation are modest, and it is therefore possible to include the approximation in existing estimation routines for nonlinear DSGE models. The most obvious estimator is probably simulated method of moments (SMM) following Duffie and Singleton (1993), as most unconditional moments can be estimated fairly accurately using samples of 1000 or 2000 observations. For instance, the Monte Carlo study in Ruge-Murcia (2012) shows that SMM performs well on simulated samples of this length. Another possibility is indirect inference with moments obtained from a vector autoregression, as suggested by Smith (1993). The simulation results in Ruge-Murcia (2007) show that this estimator also performs well when using simulated samples of 1000 or 2000 observations. Another alternative is the nonlinear filtering approach in Andreasen (2013), where the central difference Kalman filter (CDKF) computes a quasi log-likelihood function, which is maximized. This estimator is appealing because it requires relatively few evaluations of the policy function, and hence is feasible to implement with extended perturbation. To see this, consider a 30-year sample of quarterly data ($T = 120$), where a model with four states ($n_x = 4$) is estimated using four control variables as observables ($n_y = 4$). In this case, it requires only $T \times 2(n_y + n_x + 1) = 2160$ evaluations of the policy functions to compute the quasi log-likelihood function. When implementing one of these classical estimators, it is desirable to exploit multiprocessing in the optimization step to reduce the execution time. The class of evolutionary optimizers is well suited for multiprocessing, for instance, the one proposed by Andreasen (2010).

From a Bayesian perspective, nonlinear DSGE models may be estimated by a particle filter to approximate the log-likelihood function, as illustrated in An and Schorfheide (2007), Fernández-Villaverde and Rubio-Ramírez (2007), among others. However, the particle filter requires several thousand evaluations of the policy functions to approximate the log-likelihood function in a given period, implying that this approach is currently not computationally feasible with extended perturbation. Instead, Bayesian inference may be carried out using the approach in Kim (2002), where a limited-information log-likelihood function is obtained from the objective function in SMM with the optimal weighting matrix. Another alternative is to conduct Bayesian inference using an approximated likelihood function for the moment conditions as in Creel and Kristensen (2013), which also is computational feasible with extended perturbation.

[11] These computations are done on a standard desktop using Windows 10 with an Intel(R) Core(TM) i7-7600U CPU with 2.80 GHz.

6. The zero lower bound in the New Keynesian model

Following the financial crisis starting in 2008, an important requirement for the New Keynesian model is to respect the zero lower bound for the monetary policy rate. Inspired by the work of Black (1995), this is typically done by truncating the Taylor rule for $\log R_t$ at zero (see Fernandez-Villaverde, Gordon, Guerron-Quintana, and Rubio-Ramirez (2015), Gust et al. (2017), Arouba, Cuba-Borda, and Schorfheide (2018), among many others). However, such a shadow rate extension of the Taylor rule makes the New Keynesian model very demanding to solve because it introduces a kink in the model and several steady states. Higher-order perturbation methods are therefore ill-suited for solving this version of the New Keynesian model. As an alternative, Section 6.1 proposes a new Taylor rule that respects the ZLB *and* allows us to solve the New Keynesian model by higher-order perturbation methods. Section 6.2 illustrates some of the implications of applying this new Taylor rule within the New Keynesian model.

6.1 *A new ZLB consistent Taylor rule*

This section presents a Taylor rule that accounts for the ZLB without introducing the computational challenges related to a shadow rate extension of this rule. Our suggestion is simply to formulate a Taylor rule for the net interest rate r_t instead of the gross interest rate $R_t \equiv 1 + r_t$ as done in (16). That is, we propose the rule

$$\log(r_t/r_{ss}) = \rho \log(r_{t-1}/r_{ss}) + \frac{1-\rho}{r_{ss}}(\kappa_\pi \hat{\pi}_t + \kappa_y \hat{y}_t), \tag{22}$$

$\Updownarrow$

$$r_t = r_{ss} \exp\left\{\rho \log(r_{t-1}/r_{ss}) + \frac{1-\rho}{r_{ss}}(\kappa_\pi \hat{\pi}_t + \kappa_y \hat{y}_t)\right\},$$

where κ_π and κ_y are scaled by $1/r_{ss} > 0$ to make them comparable to the corresponding parameters in (16).[12] This rule obviously accounts for the ZLB without introducing a kink and is therefore suitable for higher-order perturbation methods, including extended perturbation. We also note that (22) nests the stylized rule in Benhabib, Schmitt-Grohe, and Uribe (2001) when $\rho = 0$ and $\kappa_y = 0$, and hence implies two steady-state solutions; one at the ZLB and one away from the ZLB. We follow the common approach in the literature and focus on the solution with a steady state away from the ZLB.

The first row in Table 5 reports the OLS estimates of the Taylor rule in (22) when using quarterly US data from 1959 to 2020 for the effective federal funds rate, core CPI inflation, and GDP in deviation from its potential level. We find a standard level of interest rate smoothing with $\hat{\rho} = 0.93$, and that the Federal Reserve assigns more weight to closing the inflation gap with $\hat{\kappa}_\pi = 1.62$ than the output gap with $\hat{\kappa}_y = 0.91$. These estimates are very similar to those obtained for the standard Taylor rule in row (2), which does not

[12] To see this, note for instance that $\frac{\partial r_t}{\partial \hat{\pi}} = r_t \frac{1-\rho}{r_{ss}} \kappa_\pi$ in (22), whereas the standard Taylor rule in (16) implies $\frac{\partial r_t}{\partial \hat{\pi}} = (1 + r_t)(1-\rho)\kappa_\pi$. Thus, when $r_t \approx r_{ss}$, we have that $(1 + r_t) \approx \frac{r_t}{r_{ss}}$ and the two Taylor rules give comparable partial effects.

TABLE 5. Estimated Taylor rules. The Taylor rule for the net interest rate $\log r_t = \alpha_0 + \rho \log r_{t-1} + a_\pi \hat{\pi}_t + a_y \hat{y}_t + \varepsilon_t$ is estimated in rows (1), (4), and (5). The standard Taylor rule $\log R_t = \alpha_0 + \rho \log R_{t-1} + a_\pi \hat{\pi}_t + a_y \hat{y}_t + \varepsilon_t$ with $R_t \equiv 1 + r_t$ is estimated in row (2). The shadow Taylor rule $\log R_t = \max\{0, \alpha_0 + \rho \log R_{t-1} + a_\pi \hat{\pi}_t + a_y \hat{y}_t\} + \varepsilon_t$ is estimated in row (3). In rows (2) and (3), $\hat{\kappa}_\pi = \hat{a}_\pi/(1-\hat{\rho})$ and $\hat{\kappa}_y = \hat{a}_y/(1-\hat{\rho})$, whereas the corresponding estimates are given by $\hat{a}_\pi r_{ss}/(1-\hat{\rho})$ and $\hat{a}_y r_{ss}/(1-\hat{\rho})$, respectively, in rows (1), (4), and (5) with $r_{ss} = 1/0.99 - 1$ as in specification $\mathcal{A}$. The IV estimation in row (4) is implemented using $\hat{\pi}_{t-1}$ and $\hat{y}_{t-1}$ as instruments for $\hat{\pi}_t$ and $\hat{y}_t$. The figures in parenthesis are Newey–West standard errors computed with 4 lags. Unless stated otherwise, the estimation is carried out on quarterly data from 1959 Q1 to 2020 Q1.

		ρ	a_π	a_y	κ_π	κ_y	R^2
	OLS estimation:						
(1)	New Taylor rule: $\log r_t$	0.93	11.35	6.37	1.62	0.91	0.969
		(0.03)	(5.49)	(2.51)			
(2)	Standard Taylor rule: $\log R_t$	0.91	0.11	0.04	1.12	0.38	0.952
		(0.04)	(0.07)	(0.01)			
(3)	Shadow Taylor rule	0.92	0.10	0.04	1.31	0.53	0.953
		(0.04)	(0.07)	(0.01)			
	Robustness:						
(4)	New Taylor rule: $\log r_t$ by IV	0.97	4.75	2.60	1.54	0.84	0.966
		(0.02)	(2.52)	(0.82)			
(5)	New Taylor rule: $\log r_t$ by OLS from 1985	0.80	71.19	15.00	3.66	0.77	0.976
		(0.06)	(26.16)	(4.27)			

respect the ZLB, and its shadow rate extension in row (3). The bottom part of Table 5 shows that the OLS estimates of the new Taylor rule in (22) are robust to instrumenting $\hat{\pi}_t$ and $\hat{y}_t$ by their lagged values (row 4), and that starting the sample in 1985 leads to the familiar increase in the policy response to the inflation gap ($\hat{\kappa}_\pi = 3.66$).

The column on the far right in Table 5 shows that the new Taylor rule in (22) implies an $R^2 = 0.97$, and hence provides a marginal better in-sample fit than the standard Taylor-rule and its shadow rate extension, which both have an $R^2 = 0.95$. To see where some of this improvement comes from, consider the top chart in Figure 2, which shows the effective federal funds rate and the fitted values from our new Taylor rule during the ZLB period. We first note that the policy rate never equals zero but stays about 10 to 15 basis points above the lower bound during this period. This feature is consistent with our new Taylor rule, which almost perfectly matches the policy rate at the lower bound. The bottom chart in Figure 2 reveals that the standard Taylor rule in (16) unsurprisingly predicts negative values for the policy rate. The fitted policy rates by its shadow rate extension are also too low, as they equal zero from the start of 2009 to the middle of 2014, although the policy rate is slightly positive in this period.

6.2 *Numerical illustrations*

This section explores some of the effects of using the proposed Taylor rule in the New Keynesian model from Section 4.1. That is, we replace (16) by (22), and express the policy rate in the consumption Euler equation by $1 + r_t$ instead of R_t. The model is solved

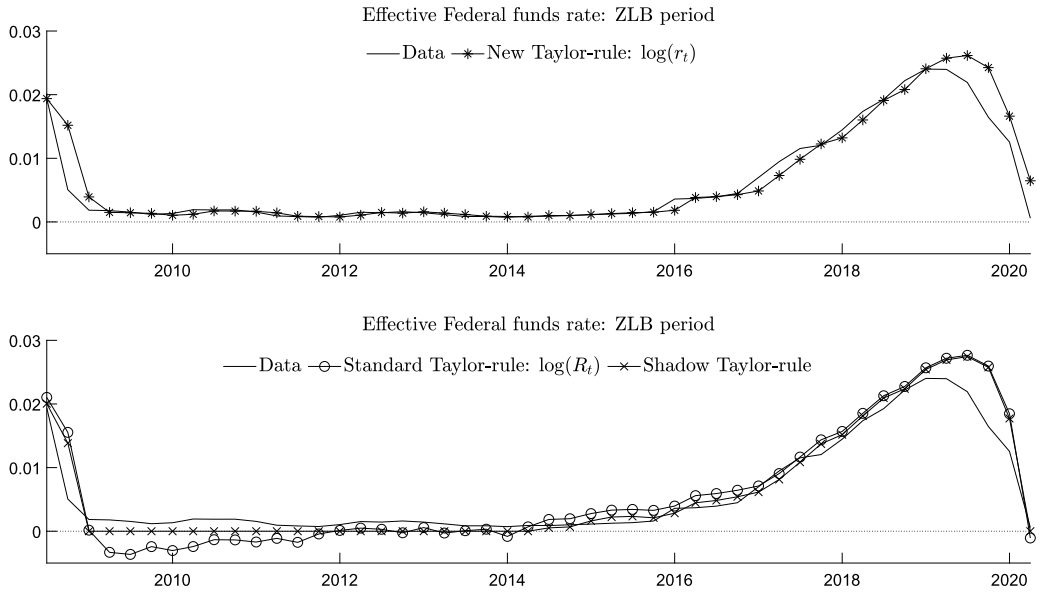

FIGURE 2. The fit of estimated Taylor rules at the ZLB. This figure shows the effective federal funds rate and the fitted values for r_t as implied by the new Taylor rule in (22), the standard Taylor rule in (16), and its shadow rate extension. These Taylor rules are estimated by OLS using quarterly data from 1959 Q1 to 2020 Q1.

by applying the standard log-transformation to all variables including r_t, as this ensures that the approximation always respects the ZLB. To simulate episodes where the economy approaches the ZLB, we omit technology shocks and follow the standard practice in the literature and introduce a preference shock d_t to the subjective discount factor β. We assume that $\log d_{t+1} = \rho_d \log d_t + \sigma_d \epsilon_{d,t+1}$, where $\epsilon_{d,t+1} \sim \mathcal{NID}(0, 1)$, $\sigma_d = 0.01$, and $\rho_d = 0.98$. For the Taylor rule, we use the estimates reported in row (1) in Table 5, while the remaining parameters attain the values in Table 2 for specification $\mathcal{A}$. We then study the performance of standard and extended perturbation, where we focus on a second order approximation, as the accuracy of standard perturbation deteriorates rapidly when increasing the approximation order and additional higher-order uncertainty terms have only a small impact on extended perturbation.

Figure 3 shows that this version of the New Keynesian model implies that the short rate approaches the ZLB in a smooth manner when lowering $\hat{d}_t$. This generates a substantial reduction in inflation $\hat{\pi}_t$, which is much larger than implied by a linear solution that does not respect the ZLB. This linear solution is obtained by a first-order approximation to the solution from extended perturbation around $\hat{d}_t = 0$. At the ZLB, we find that consumption falls by more with extended perturbation than implied by this linear solution. Thus, the implications of enforcing the ZLB by the proposed Taylor rule in (22) are very similar to those reported for a shadow rate extension of the standard Taylor rule (see, for instance, Fernandez-Villaverde et al. (2015)). The bottom right chart in Figure 3 reveals that extended perturbation provides a fairly accurate approximation at the ZLB, and that it generates lower Euler errors when compared to standard perturbation.

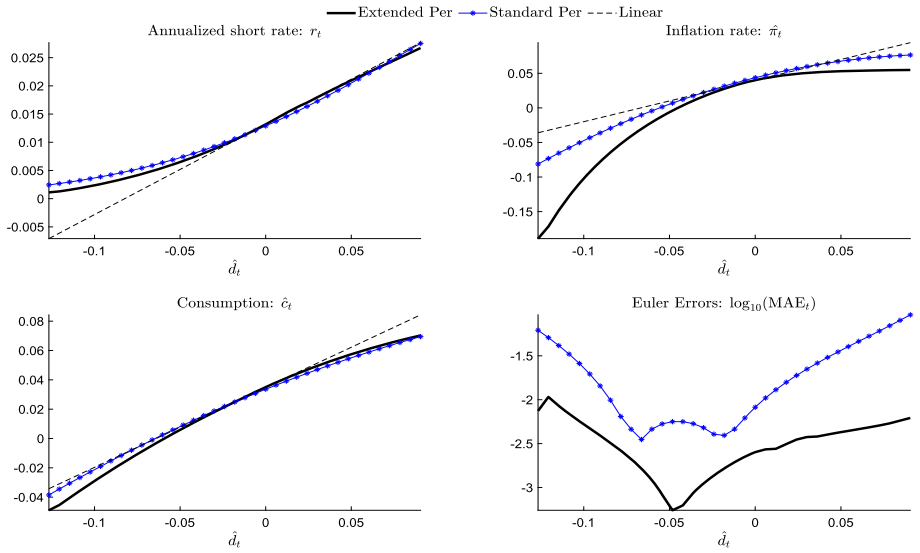

FIGURE 3. Policy functions at the ZLB with the new Taylor rule. This figure shows the level of the annualized policy rate r_t, inflation $\hat{\pi}_t$, consumption $\hat{c}_t$, and the Euler-equation errors $\log_{10}(\text{MAE}_t)$ as a function of the preference shock $\hat{d}_t$ in the New Keynesian model with the new Taylor rule in (22) when solved by standard and extended perturbation to second order. The remaining states are fixed at $\hat{k}_t = -2\sigma_k$, $\hat{r}_{t-1} = -2.5\sigma_r$, and $\hat{s}_t = 0$, where σ_k and σ_r denote the standard deviation of $\hat{k}_t$ and $\hat{r}_{t-1}$, respectively, in a log-linear approximation. The reported linear solution is obtained by a first-order approximation to the solution from extended perturbation around $\hat{d}_t = 0$.

The top chart in Figure 4 shows a simulated sample of 1000 observations for the annualized short rate. We find that our new Taylor rule frequently takes the monetary policy rate close to the ZLB. From observations 875 to 936, the short rate is below 50 basis points, meaning that the economy is close to the ZLB for 61 consecutive quarters, or slightly more than 15 years. This shows that the new Taylor rule in (22) can generate long stays at the ZLB. The bottom chart in Figure 4 reveals that extended perturbation delivers a reasonable degree of accuracy with $\log_{10}\text{MAE}_t$ (for the entire model) being around -3, also at the ZLB. In comparison, the corresponding series for the short rate by standard perturbation follows the one from extended perturbation fairly closely until observation 690, where standard perturbation explodes.

7. CONCLUSION

This paper introduces the extended perturbation approximation, which improves the performance of standard perturbation by using a more accurate solution for the certainty equivalent component of the policy function. For the New Keynesian model with Calvo pricing, we show that extended perturbation achieves higher accuracy than standard perturbation. We also show that the gain in accuracy is sufficiently large to generate a stable approximation by extended perturbation when standard perturbation explodes.

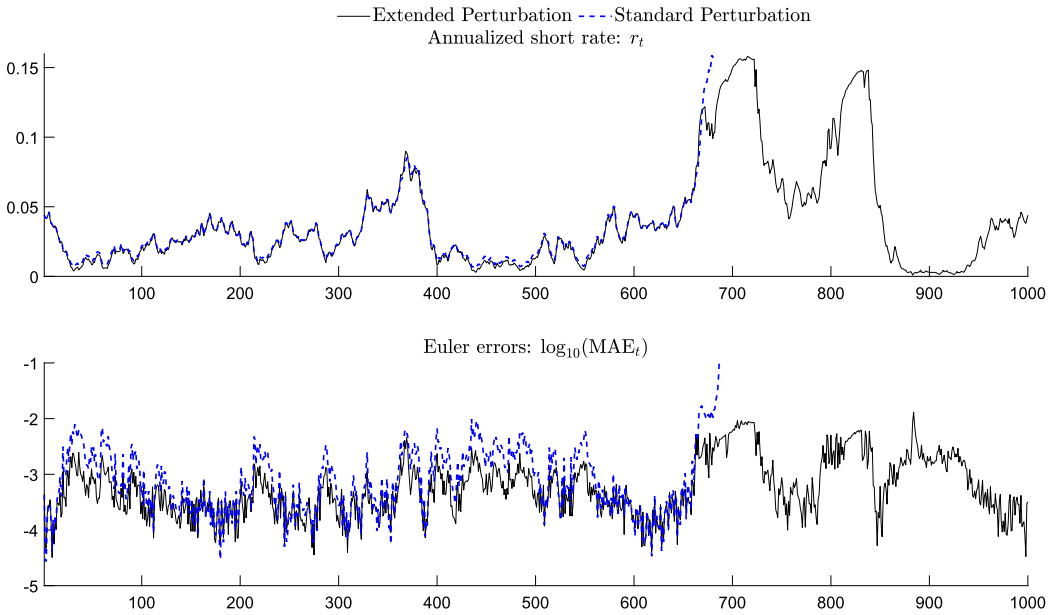

FIGURE 4. Simulations from the New Keynesian model enforcing the ZLB. The top chart shows the level of the annualized policy rate r_t in a simulated sample of 1.000 observations (with a burn-in of 100 observations) in the New Keynesian model with the new Taylor rule in (22) when solved by standard and extended perturbation to second order. The bottom chart shows the Euler equation errors $\log_{10}(\text{MAE}_t)$ for the entire model at a given point in the simulation.

Our results therefore suggest that the explosive behavior of standard perturbation reported in the literature may be related to inaccuracies in the certainty equivalent component of the policy function that may be eliminated by using extended perturbation. To reduce the computational costs of extended perturbation, we also introduce several improvements of the extended path, which substantially lower execution costs and make extended perturbation feasible for estimation. In an application, we use extended perturbation on a New Keynesian model that enforces the ZLB by considering a Taylor rule for the log-transformed net interest rate instead of the widely used shadow rate specification. The results show that this modified version of the New Keynesian model can generate long stays at the ZLB and is easy to solve using extended perturbation, which provides a very accurate solution.

APPENDIX: UNCERTAINTY ADJUSTMENT BY STANDARD PERTURBATION

We clearly have that the stochastic part of the policy function is zero when we omit uncertainty ($\sigma = 0$), that is, $\mathbf{g}^{\text{stoch}}(\mathbf{x}_t, \sigma = 0) = \mathbf{0}$ and $\mathbf{h}^{\text{stoch}}(\mathbf{x}_t, \sigma = 0) = \mathbf{0}$ for all values of $\mathbf{x}_t$. This implies that all derivatives of $\mathbf{g}$ and $\mathbf{g}^{\text{CE}}$ solely with respect to the state variables are identical at $\sigma = 0$, and similarly for $\mathbf{h}$ and $\mathbf{h}^{\text{CE}}$. That is,

$$\begin{aligned} \mathbf{g}(\mathbf{x}_t, \sigma = 0)_{\mathbf{x}^m} &= \mathbf{g}^{\text{CE}}(\mathbf{x}_t)_{\mathbf{x}^m} \quad \text{for all } \mathbf{x}_t \in \mathcal{X}_x, \\ \mathbf{h}(\mathbf{x}_t, \sigma = 0)_{\mathbf{x}^m} &= \mathbf{h}^{\text{CE}}(\mathbf{x}_t)_{\mathbf{x}^m} \quad \text{for all } \mathbf{x}_t \in \mathcal{X}_x \end{aligned} \tag{23}$$

for $m = \{0, 1, 2, \ldots\}$, where subscripts refer to partial derivatives taken m times with respect to $\mathbf{x}_t$. We also note that all derivatives involving the perturbation parameter σ are identical for $\mathbf{g}$ and $\mathbf{g}^{\text{stoch}}$ because σ does not appear in $\mathbf{g}^{\text{CE}}$, and similarly for $\mathbf{h}$ and $\mathbf{h}^{\text{stoch}}$. That is,

$$\begin{aligned} \mathbf{g}(\mathbf{x}_t, \sigma)_{\mathbf{x}^m\sigma^j} &= \mathbf{g}^{\text{stoch}}(\mathbf{x}_t, \sigma)_{\mathbf{x}^m\sigma^j} \quad \text{for all } \mathbf{x}_t \in \mathcal{X}_x, \sigma \in \mathbb{R}_+, \\ \mathbf{h}(\mathbf{x}_t, \sigma)_{\mathbf{x}^m\sigma^j} &= \mathbf{h}^{\text{stoch}}(\mathbf{x}_t, \sigma)_{\mathbf{x}^m\sigma^j} \quad \text{for all } \mathbf{x}_t \in \mathcal{X}_x, \sigma \in \mathbb{R}_+ \end{aligned} \tag{24}$$

for $m = \{0, 1, 2, \ldots\}$ and $j = \{1, 2, \ldots\}$, where subscripts refer to partial derivatives taken m times with respect to $\mathbf{x}_t$ and j times with respect to σ. The standard perturbation method exploits partial derivatives of the model in (1) to first compute partial derivatives of the $\mathbf{g}$- and $\mathbf{h}$-functions at the steady state with respect to $\mathbf{x}_t$. These partial derivatives are subsequently used to compute partial derivatives of the $\mathbf{g}$- and $\mathbf{h}$-functions at the steady state with respect to σ and $\mathbf{x}_t$. But (23) and (24) show that these partial derivatives are unaffected by approximating $\mathbf{g}^{\text{CE}}$ and $\mathbf{h}^{\text{CE}}$ by the extended path instead of the standard Taylor series expansion. As a result, the standard perturbation method can be used to compute $\mathbf{g}^{\text{stoch}}(\mathbf{x}_t, \sigma)_{\mathbf{x}^m\sigma^j}$ and $\mathbf{h}^{\text{stoch}}(\mathbf{x}_t, \sigma)_{\mathbf{x}^m\sigma^j}$ at the steady state when the certainty equivalent part of the policy function is computed by the extended path.

References

Adjemian, Stephane and Michel Juillard (2010), "Dealing with ZLB in DSGE models: An application to the Japanese economy." ESRI Discussion Paper No. 258. [1191]

Adjemian, Stephane and Michel Juillard (2013), "Stochastic extended path approach." Working Paper. [1178, 1179]

Ajevskis, Viktors (2017), "Semi-global solutions to DSGE models: Perturbation around a deterministic path." *Studies in Nonlinear Dynamics and Econometrics*, 21 (2), 1–28. [1179, 1180]

An, Sungbau and Frank Schorfheide (2007), "Bayesian analysis of DSGE models." *Econometric Review*, 26 (2–4), 113–172. [1171, 1193]

Andreasen, Martin M. (2010), "How to maximize the likelihood function for a DSGE model." *Computational Economics*, 35 (2), 127–154. [1193]

Andreasen, Martin M. (2012), "On the effects of rare disasters and uncertainty shocks for risk premia in non-linear DSGE models." *Review of Economic Dynamics*, 15 (3), 295–316. [1174, 1176]

Andreasen, Martin M. (2013), "Non-linear DSGE models and the central difference Kalman filter." *Journal of Applied Econometrics*, 28, 929–955. [1193]

Andreasen, Martin M., Jesus Fernandez-Villaverde, and Juan F. Rubio-Ramirez (2018), "The pruned state space system for non-linear DSGE models: Theory and empirical applications to estimation." *Review of Economic Studies*, 85 (1), 1–49. [1172]

Andreasen, Martin M., Anders F. Kronborg (2022), "Supplement to 'The extended perturbation method: With applications to the New Keynesian model and the zero lower bound'." *Quantitative Economics Supplemental Material*, 13, https://doi.org/10.3982/QE1102.

Arouba, S. Boragan, Pablo Cuba-Borda, and Frank Schorfheide (2018), "Macroeconomic dynamics near the ZLB: A tale of two countries." *Review of Economic Studies*, 85 (1), 87–118. [1173, 1194]

Beaudry, Paul, Dana Galizia, and Franck Portier (2020), "Putting the cycle back into the business cycle analysis." *American Economic Review*, 110 (1), 1–47. [1174]

Benhabib, Jess, Stephanie Schmitt-Grohe, and Martin Uribe (2001), "The perils of Taylor rules." *Journal of Economic Theory*, 96, 40–69. [1194]

Black, Fischer (1995), "Interest rates as options." *The Journal of Finance*, 50 (5), 1371–1376. [1194]

Blasques, Francisco and Marc Nientker (2020), "Transformed perturbation solutions for dynamic stochastic general equilibrium models," 1–57. Working Paper. [1183]

Boucekkine, Raouf (1995), "An alternative methodology for solving nonlinear forward looking models." *Journal of Economic Dynamic and Control*, 19 (4), 711–734. [1181, 1190]

Burnside, Craig (1998), "Solving asset pricing models with Gaussian shocks." *Journal of Economic Dynamics and Control*, 22, 329–340. [1177]

Calvo, Guillermo A. (1983), "Staggered prices in a utility-maximizing framework." *Journal of Monetary Economics*, 12 (3), 383–398. [1183, 1184]

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans (2005), "Nominal rigidities and the dynamic effects of a shock to monetary policy." *Journal of Political Economy*, 113 (1), 1–45. [1189]

Creel, Michael and Dennis Kristensen (2013), "Indirect likelihood inference." BSE Working Paper. [1193]

Duffie, Darrell and Kenneth J. Singleton (1993), "Simulated moments estimation of Markov models of asset prices." *Econometrica*, 61 (4), 929–952. [1193]

Fair, Ray C. and John B. Taylor (1983), "Solution and maximum likelihood estimation of dynamic nonlinear rational expectations models." *Econometrica*, 51 (4), 1169–1185. [1172, 1178]

Fernandez-Villaverde, Jesus, Grey Gordon, Pablo Guerron-Quintana, and Juan F. Rubio-Ramirez (2015), "Nonlinear adventures at the zero lower bound." *Journal of Economic Dynamic and Control*, 57, 182–204. [1194, 1196]

Fernández-Villaverde, Jesús, Pablo Guerrón-Quintana, Juan F. Rubio-Ramírez, and Martin Uribe (2011), "Risk matters: The real effects of volatility shocks." *American Economic Review*, 101 (6), 2530–2561. [1171]

Fernández-Villaverde, Jesús and Juan F. Rubio-Ramírez (2007), "Estimating macroeconomic models: A likelihood approach." *Review of Economic Studies*, 74 (4), 1–46. [1193]

Gust, Christopher, Edward Herbst, David Lopez-Salido, and Matthew E. Smith (2017), "The empirical implications of the interate-rate lower bound." *American Economic Review*, 107 (7), 1971–2006. [1173, 1194]

Haan, Wouter J. Den and Joris De Wind (2012), "Nonlinear and stable perturbation-based approximations." *Journal of Economic Dynamic and Control*, 36 (10), 1477–1497. [1172, 1183]

Jermann, Urban J. (1998), "Asset pricing in production economics." *Journal of Monetary Economics*, 41 (2), 257–275. [1184]

Judd, Kenneth L. and Sy-Ming Guu (1997), "Asymptotic methods for aggregate growth models." *Journal of Economic Dynamics and Control*, 21 (6), 1025–1042. [1175]

Juillard, Michel and Sebastien Villemot (2011), "Multi-country real business cycle models: Accuracy tests and test bench." *Journal of Economic Dynamic and Control*, 35, 178–185. [1178]

Kim, Jae-Young (2002), "Limited information likelihood and Bayesian analysis." *Journal of Econometrics*, 107, 175–193. [1193]

Kim, Jinill, Sunghyun Kim, Ernst Schaumburg, and Chrisopher A. Sims (2008), "Calculating and using second-order accurate solutions of discrete time dynamic equilibrium models." *Journal of Economic Dynamics and Control*, 32 (11), 3397–3414. [1172]

Kim, Jinill and Francisco J. Ruge-Murcia (2009), "How much inflation is necessary to grease the wheels?" *Journal of Monetary Economics*, 56 (3), 365–377. [1171]

King, Rober G. and Sergio T. Rebelo (1999), "Resuscitating real business cycles." *Handbook of Macroeconomics*, 1, 927–1007. [1174]

Levintal, Oren (2017), "Fifth-order perturbation solution to DSGE models." *Journal of Economics Dynamic and Control*, 80, 1–16. [1179]

Lombardo, Giovanni and Harald Uhlig (2018), "A theory of pruning." *International Economic Review*, 59 (4), 1825–1836. [1172]

Lubik, Tohomas A. and Frank Schorfheide (2004), "Testing for indeterminancy: An application to U.S. monetary policy." *American Economic Review*, 94 (1), 190–217. [1174]

Lucas, Robert E. (1978), "Asset prices in an exchange economy." *Econometrica*, 46 (6), 1429–1445. [1177]

Pötscher, Benedikt M. and Ingmar Prucha (1997), *Dynamic Nonlinear Econometric Models: Asymptotic Theory*. Springer-Verlag, Berlin Heidelberg, New York. [1180]

Rietz, Thomas A. (1988), "The equity premium: A solution." *Journal of Monetary Economics*, 22, 117–131. [1176]

Rudebusch, Glenn D. and Eric T. Swanson (2012), "The bond premium in a DSGE model with long-run real and nominal risks." *American Economic Journal: Macroeconomics*, 4 (1), 1–43. [1171]

Ruge-Murcia, Francisco (2012), "Estimating nonlinear DSGE models by the simulated method of moments: With an application to business cycles." *Journal of Economics Dynamic and Control*, 36 (6), 914–938. [1193]

Ruge-Murcia, Francisco J. (2007), "Methods to estimate dynamic stochastic general equilibrium models." *Journal of Economic Dynamics and Control*, 31 (8), 2599–2636. [1193]

Schmitt-Grohé, Stephanie and Martin Uribe (2004), "Solving dynamic general equilibrium models using a second-order approximation to the policy function." *Journal of Economic Dynamics and Control*, 28 (4), 755–775. [1176]

Smith, Anthony A. Jr. (1993), "Estimating nonlinear time-series models using simulated vector autoregressions." *Journal of Applied Econometrics*, 8 (Supplement: Special Issue on Econometric Inference Using Simulation Techniques), S63–S84. [1193]

Co-editor Tao Zha handled this manuscript.

Manuscript received 23 March, 2018; final version accepted 21 November, 2021; available online 9 January, 2022.