Red Nacional de Investigadores en Economía

Piecewise Linear Solutions for Non-Stationary Models

Mariano Kulish (University of Sydney)

Inna Tsener (Universitat de les Illes Balears)

DOCUMENTO DE TRABAJO N° 387

Febrero de 2026

Citar como:
Kulish, Mariano, Inna Tsener (2026). Piecewise Linear Solutions for Non-Stationary Models. Documento de trabajo RedNIE N°387.

Piecewise Linear Solutions for Non-Stationary Models[*]

Mariano Kulish[†] and Inna Tsener[‡]

3 October 2025

Abstract

We assess the accuracy and efficiency of time-varying linear solution methods for non-stationary rational expectations models. These methods construct a sequence of local linear approximations, each with coefficients that vary over time, based on a set of expansion points. Benchmarking against globally accurate non-linear solutions, we show, both theoretically and numerically, that their accuracy depends critically on the choice of expansion points and on agents' expectations about the future. Our results contribute to the literature on solving non-stationary stochastic models with rational expectations, spanning a wide range of sources of non-stationarity, including evolving structural parameters, changing policy regimes, and cases without a balanced growth path.

JEL Codes: C61, C63, C68

Keywords: piecewise linear solutions, approximation points, time-inhomogeneous models, non-stationary models, semi-Markov models, unbalanced growth, time-varying parameters, extended function path

---

[*]We thank Fernando Alvarez, Marco Bassetto, Ippei Fujiwara, Alexander Meyer-Gohde, Simon Mongey and Juan Pablo Nicolini for useful discussions. We also thank seminar and conference participants at the Hitotsubashi University, QuantEcon, Keio University, Federal Reserve Bank of Minneapolis, Federal Reserve Board, EDMM and CEM at the University of Adelaide, UTDT, Universitat de Barcelona, Sydney Macro Reading Group, Computing in Economics and Finance 2024, CFE-CMStatistics and CESC. Financial support of the Spanish Ministry of Education is gratefully acknowledged, grant PID2022-138428NA-I00 funded by MCIN/AEI/10.13039/501100011033 and by "ERDF A way of making Europe".

[†]School of Economics, University of Sydney; mariano.kulish@sydney.edu.au.

[‡]Department of Applied Economics, Universitat de les Illes Balears.

1 Introduction

Conventional solution methods for dynamic models with rational expectations rely on a stationary environment. In models with growing productivity, for instance, it is standard to normalize variables so that the system becomes stationary and conventional methods can be applied. After such a transformation, the parameters governing preferences, technologies, and policy rules are assumed to remain constant.

Our interest lies in situations where the assumption of a stationary environment does not hold, or where no stationary representation is available. The situations we have in mind are not abstract possibilities but real-world cases of practical importance. Examples include economies without a balanced growth path, as would be the case for models of structural transformation; economies transitioning from one balanced growth path to another; or economies experiencing disruptive technical change, such as the automation and AI revolutions. Other examples include changes in policy regimes – for instance, when joining the euro[1] – as well as permanent disinflations, fiscal reforms, or the zero lower bound constraint on policy rates.

This paper builds on the time-varying linear solution method introduced in Cagliarini and Kulish (2013) and Kulish and Pagan (2017), which provides a tractable approach to solving rational expectations models with evolving parameters. That earlier work focused on models already expressed in linearized form and did not assess how well the solution approximated the underlying non-linear dynamics.[2]

In contrast, we begin with the fully non-linear model and treat the piecewise linear solution as an approximation method.[3] Our goal is to evaluate the accuracy and efficiency of this approximation in non-stationary settings, including cases where the steady state shifts over time or where no steady state exists.

We find that two modeling choices – (i) how expansion points are constructed, and (ii)

[1] See Gomez-Gonzalez and Rees (2018) for an analysis that accounts for a break in the conduct of monetary policy when Spain joined the euro.

[2] Jones (2017) and Guerrieri and Iacoviello (2015) compare a similar piecewise linear solution to a non-linear solution but do this for the case of a zero lower bound occassionally binding constraint for which the steady state of the economy under the reference regime stays constant.

[3] We use the term 'piecewise linear' to refer to a solution that is linear at each point in time, with coefficients that vary over time. Some researchers prefer the term time-varying coefficients to avoid confusion with a solution that is piecewise linear in the state space. However, the terminology 'piecewise linear' is also common in applications of the time-varying VAR approach to occasionally binding constraints, where coefficients vary with the realization of shocks and the duration of the binding constraint. .

how expectations are modeled (anticipated vs unanticipated changes) – have first-order effects on accuracy. These are not mere implementation details; our results show they are central to the quality of the approximation.

By benchmarking against global non-linear solutions, we validate the method in a non-linear context. Moreover, because it retains the tractability of a linear structure, the piecewise linear solution is applicable to large scale models, including those without a balanced growth path, changing policy regimes, or structural drifts, extending its use far beyond previous applications.

Related to this paper, Maliar et al. (2020) propose a non-linear solution method – known as the extended function path (EFP) – for non-stationary models. While their approach is accurate, it is computationally intensive and often infeasible for medium- or large-scale applications. In particular, estimation or policy exercises typically require solving the model many times, making global methods prohibitively slow. In contrast, the piecewise linear solution, by preserving a linear structure at each point in time, scales well with model size and remains relatively computationally efficient. These features make it a promising candidate for applied work – provided we understand the conditions under which it delivers a good approximation.

To uncover these conditions, we use a simple stochastic growth model as a laboratory. Rather than working with detrended variables, we operate in levels, thereby preserving the non-stationary structure of the model. We consider two forms of non-stationarity: one driven by trend labour-augementing productivity growth, and another driven by changes in capital intensity. The first case is useful because it allows a direct comparison between the piecewise linear solution in levels and the global non-linear solution expressed in detrended variables. The second case goes further – non-stationarity arises in a way that precludes a stationary transformation, allowing us to benchmark the piecewise linear solution against the EFP method.

Following Maliar et al. (2020), we assume that the non-stationarity eventually stops at some distant future date. From that point onward, the model becomes stationary and can be solved using standard methods. This effectively transforms the problem into one with a finite horizon of non-stationarity, allowing us to solve backward recursively from a known terminal regime, similar to lifecycle models.

The accuracy of the piecewise linear approximation hinges on the conditions of the Turnpike Theorem, which states that the optimal path of a finite-horizon economy remains close to its infinite-horizon counterpart for most of the horizon, regardless of initial or terminal conditions.[4] When these conditions hold – as they do in the stochastic growth model – the piecewise linear approximation closely tracks the infinite-horizon non-linear global solution over the periods of interest.

Our main finding is that the piecewise linear solution is most accurate when expansion points evolve in line with the non-stochastic path the economy would have followed in the absence of shocks, and when agents anticipate the structural evolution of the economy. In the appendix, we formally show why the non-stochastic path provides the right basis for linearization: if expansion points deviate from this path, the approximation incurs a first-order bias. Moreover, in the case of the stochastic growth model, agent anticipation of the path of labour-augmenting productivity ensures consistency with the model's information structure.

Although we apply the piecewise linear solution to deal with a stochastic dynamic forward looking model, the intuition for our main finding can be grasped with a simple example. Figure 1 shows that one can approximate any highly non-linear function by using many linear functions. The approximation will be more accurate when more linear functions and expansion points are used, assuming these are 'good' expansion points.

Figure 1: Piecewise Linear Approximation

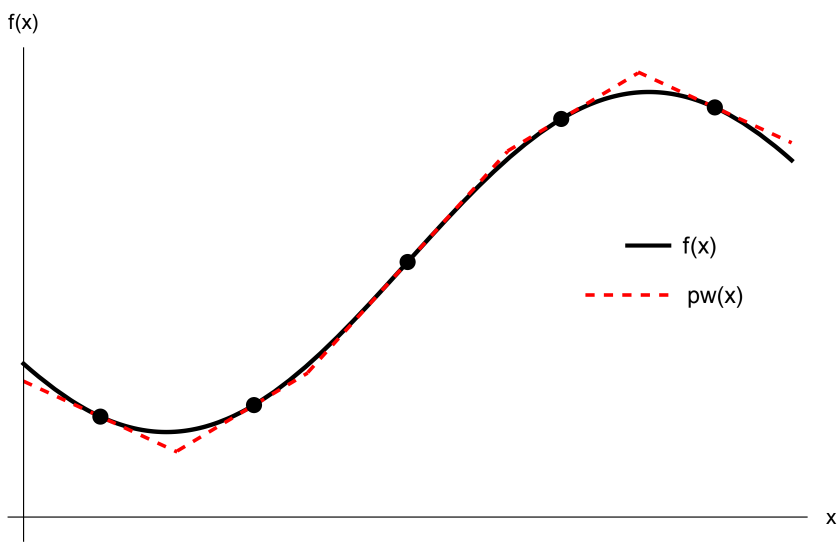

This paper is related to two major lines of research that study nonstationary models similar to the one considered in this paper. Firstly, similar to Maliar et al. (2020) we study

[4] See McKenzie (1976) for a discussion of turnpike theory.

solutions that take a form of a sequence of time-varying policy functions. Specifically, using global approximation methods Maliar et al. (2020) construct a path of time-inhomogeneous Markov decision functions. Secondly, Evans and Phillips (2015) and Phillips (2017), propose a method for solving non-stationary rational expectations models using local approximation techniques. These papers, however, rely on the simplifying assumption of linearizing around the current state to construct the sequence of time-varying matrices. By evaluating alternative expansion paths, we show that approaches that do not rely on the non-stochastic path from the non-linear model as the expansion basis yield inaccurate solutions in non-stationary contexts. Similar to Mennuni et al. (2025) and Ajevskis (2017), we rely on a dynamic first-order Taylor expansion. However, there are important differences. Mennuni et al. (2025) constructs an equilibrium path using different first-order approximations at different points in the state space. Ajevskis (2017) works with stationary models that have a well-defined steady state, but applies his method starting from arbitrary initial conditions away from steady state. In contrast, we apply time-varying expansion points to genuinely non-stationary models, and rely on the Turnpike Theorem to obtain an accurate solution over a long but finite horizon.

The rest of the paper is structured as follows. Section 2 presents a version of the stochastic growth model that we use to assess the performance of piecewise linear solution methods. This model is deliberately kept simple to allow comparison with global non-linear solutions, which are otherwise computationally infeasible in more complex settings. For completeness, Section 3 reviews the piecewise linear solution proposed by Kulish and Pagan (2017). Section 4 discusses possible expansion points that researchers may consider for their non-linear models. Section 5 reports numerical results, and Section 6 concludes.

2 Non-Stationary Neoclassical Growth Models

To illustrate the piecewise linear approach to solving non-stationary models, we consider a stochastic neoclassical growth model in which a representative agent inelastically supplies one unit of labour and solves:

$$\max_{\{C_t,K_t\}_{t=0}^{\infty}} \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t \frac{C_t^{1-\gamma} - 1}{1-\gamma} \right] \tag{1}$$

$$\text{s.t.} \quad C_t + K_t = (1-\delta)K_{t-1} + f_t(K_{t-1}, Z_t) \tag{2}$$

$$\log Z_t = \rho \log Z_{t-1} + \varepsilon_{Z,t} \tag{3}$$

Here, $C_t \geq 0$ and $K_t \geq 0$ denote consumption and physical capital, respectively. The parameter $\beta \in (0,1)$ is the discount factor, $\delta \in [0,1]$ is the depreciation rate, and $\varepsilon_{Z,t}$ is an i.i.d. shock. The operator $\mathbb{E}_t[\cdot]$ denotes expectations conditional on information available at time t, and the initial conditions (K_{-1}, Z_{-1}) are given.

The production function $f_t : \mathbb{R}_+^2 \to \mathbb{R}_+$ is time-dependent, making the system non-stationary.[5] We consider two sources of non-stationarity in what follows.

In the first case – standard in the literature – labour-augmenting productivity grows over time:

$$f_t(K_{t-1}, Z_t) = Z_t K_{t-1}^{\alpha} A_t^{1-\alpha}, \quad \text{with} \quad A_t = g A_{t-1}, \quad g > 1. \tag{4}$$

In the second case, the non-stationarity arises from a time-varying capital share:

$$f_t(K_{t-1}, Z_t) = Z_t K_{t-1}^{\alpha_t}, \quad \text{with} \quad \alpha_t = \alpha_0 + m_\alpha t, \quad t = 0, 1, \ldots, T. \tag{5}$$

An important distinction between the first case, where f_t follows equation (4), and the second case, where f_t follows equation (5), is that the former admits a balanced growth path and the model can be transformed into a stationary system through normalization. In contrast, the second case does not allow such a re-normalization: the presence of a drifting capital share, α_t, prevents the system from being made stationary through standard transformations.

The first case, $f_t(K_{t-1}, Z_t) = Z_t K_{t-1}^{\alpha} A_t^{1-\alpha}$, is particularly useful because it allows us to compare the piecewise linear solution applied to the model in levels with a globally accurate non-linear solution computed in normalized variables. Defining the real interest rate, R_t, as the marginal product of capital, the first-order conditions of the non-stationary model are:

$$C_t^{-\gamma} = \beta \mathbb{E}_t \left(C_{t+1}^{-\gamma} R_{t+1} \right) \tag{6}$$

$$R_t = 1 - \delta + \alpha Z_t K_{t-1}^{\alpha-1} A_t^{1-\alpha} \tag{7}$$

$$C_t + K_t = (1-\delta)K_{t-1} + Z_t K_{t-1}^{\alpha} A_t^{1-\alpha} \tag{8}$$

[5] One could, of course, introduce non-stationarity through other channels such as a time-dependent utility function or a time-varying stochastic process. For the purposes of this paper, it suffices to illustrate the approach using the production function.

Defining lower-case variables as $c_t = C_t/A_t$, $k_t = K_t/A_t$, $r_t = R_t/g$, and $Z_t = z_t$, the first-order conditions of the stationary representation are:

$$c_t^{-\gamma} = \beta g^{1-\gamma}\mathbb{E}_t\left(c_{t+1}^{-\gamma} r_{t+1}\right) \tag{9}$$

$$r_t = (1-\delta)\frac{1}{g} + \alpha z_t\left(\frac{k_{t-1}}{g}\right)^{\alpha-1}\cdot\frac{1}{g} \tag{10}$$

$$c_t + k_t = (1-\delta)\frac{k_{t-1}}{g} + z_t\left(\frac{k_{t-1}}{g}\right)^{\alpha} \tag{11}$$

To illustrate our approach, we apply the piecewise linear solution to the non-stationary system given by equations (6)–(8), where productivity grows over time, $A_t = gA_{t-1}$ with $g > 1$. We then compare it to a global non-linear solution obtained using a Smolyak projection method, following Krueger and Kubler (2004), applied to the normalized system (9)–(11).

In the second case, where $f_t(K_{t-1}, Z_t) = Z_t K_{t-1}^{\alpha_t}$ and α_t drifts over time, no stationary transformation is available. In this setting, we benchmark the piecewise linear solution against the Extended Function Path (EFP) method of Maliar et al. (2020).

To solve a non-stationary system such as (6)–(8), we assume that the source of non-stationarity stops at some distant future date T: in the first case, this means productivity growth halts; in the second, that the drift in the capital share ceases. We use the solution to the resulting stationary system from $t \geq T$ as a terminal condition and solve the model backward from that point. As discussed in the introduction, when the conditions of the Turnpike Theorem hold, the piecewise linear solution can provide an accurate approximation over most of the horizon – provided that expansion points are appropriately chosen, as we discuss below.

The system (6)–(8) is non-stationary because its optimal policy and value functions depend not only on the state but also on time. In general, non-stationary models cannot be solved using conventional numerical methods, since their solution is not a single policy function but a sequence of time-varying policy functions that satisfy the model's equilibrium conditions. Some non-stationary models admit a balanced growth path and can be transformed into stationary form, as discussed in King et al. (1988). However, such transformations are not always possible – for instance, in the case where the capital share drifts over time.

Having presented the model and the two cases of non-stationarity, we now turn to a brief review of the time-varying VAR solution proposed by Kulish and Pagan (2017).

3 The Piecewise Linear Solution

A standard approach to solving non-linear stochastic models with rational expectations, well explained in Uhlig (1999), is to first compute the non-stochastic steady state of the model, then perform a first-order Taylor expansion around that steady state. This yields a linearized system of structural equations that can be written in matrix form as

$$\mathbf{A}y_t = \mathbf{C} + \mathbf{B}y_{t-1} + \mathbf{D}\mathbb{E}_t y_{t+1} + \mathbf{F}\varepsilon_t, \tag{12}$$

where y_t is an $n \times 1$ vector of state and jump variables, and ε_t is an $l \times 1$ vector of exogenous shocks.[6] With no loss of generality, we take the shocks to be white noise. The matrices in (12) are populated by coefficients that are typically non-linear functions of the primitive parameters of the model.[7]

Kulish and Pagan (2017) take the linearized system (12) as a starting point and allow the coefficients of the structural matrices to vary over time. That is, they postulate that the linearized structural equations follow

$$\mathbf{A}_t y_t = \mathbf{C}_t + \mathbf{B}_t y_{t-1} + \mathbf{D}_t\mathbb{E}_t y_{t+1} + \mathbf{F}_t\varepsilon_t, \tag{13}$$

where the matrices $\mathbf{A}_t, \mathbf{B}_t, \mathbf{C}_t, \mathbf{D}_t, \mathbf{F}_t$ evolve with time to capture structural changes in the model. The linear rational expectations model in (13) nests the stationary case as a special case. When the model is stationary, the matrices $\mathbf{A}_t \equiv \mathbf{A}$, $\mathbf{B}_t \equiv \mathbf{B}$, $\mathbf{D}_t \equiv \mathbf{D}$, $\mathbf{C}_t \equiv \mathbf{C}$, and $\mathbf{F}_t \equiv \mathbf{F}$ are time-invariant, and the constant vector $\mathbf{C}$ plays a key role in determining the steady state of y_t.

Assume that after $t \geq T + 1$ the economy enters a terminal regime given by

$$\bar{\mathbf{A}}y_t = \bar{\mathbf{C}} + \bar{\mathbf{B}}y_{t-1} + \bar{\mathbf{D}}\mathbb{E}_t y_{t+1} + \bar{\mathbf{F}}\varepsilon_t, \tag{14}$$

with unique solution $y_t = \bar{\mathbf{J}} + \bar{\mathbf{Q}}y_{t-1} + \bar{\mathbf{G}}\varepsilon_t$. Then, as shown in Kulish and Pagan (2017), if a solution to (13) exists and is unique, it takes the form of a time-varying VAR:

$$y_t = \mathbf{J}_t + \mathbf{Q}_t y_{t-1} + \mathbf{G}_t\varepsilon_t. \tag{15}$$

[6] The system can be generalized to allow for additional lags of y_t and for expectations at different horizons or from earlier dates.

[7] The constant-parameter case in (12) can be solved using standard methods, including those of Anderson (1997), Blanchard and Kahn (1980), Binder and Pesaran (1995), King and Watson (1998), Klein (2000), Sims (2001), and Uhlig (1999).

If (15) is a solution, it implies that $\mathbb{E}_t y_{t+1} = \mathbf{J}_{t+1} + \mathbf{Q}_{t+1} y_t$. Substituting this into (13), the method of undetermined coefficients yields the following recursions for the reduced-form matrices:

$$\mathbf{J}_t = (\mathbf{I} - \Lambda_t \mathbf{Q}_{t+1})^{-1} (\mathbf{\Gamma}_t + \Lambda_t \mathbf{J}_{t+1}), \tag{16}$$

$$\mathbf{Q}_t = (\mathbf{I} - \Lambda_t \mathbf{Q}_{t+1})^{-1} \Phi_t, \tag{17}$$

$$\mathbf{G}_t = (\mathbf{I} - \Lambda_t \mathbf{Q}_{t+1})^{-1} \Psi_t, \tag{18}$$

where

$$\Phi_t \equiv \mathbf{A}_t^{-1} \mathbf{B}_t, \quad \Lambda_t \equiv \mathbf{A}_t^{-1} \mathbf{D}_t, \quad \Psi_t \equiv \mathbf{A}_t^{-1} \mathbf{F}_t, \quad \mathbf{\Gamma}_t \equiv \mathbf{A}_t^{-1} \mathbf{C}_t.$$

Starting from the terminal regime, with $\mathbf{J}_{T+1} = \bar{\mathbf{J}}$, $\mathbf{Q}_{T+1} = \bar{\mathbf{Q}}$, and $\mathbf{G}_{T+1} = \bar{\mathbf{G}}$, the recursions (16)–(18) generate the sequence of policy and decision rules, that is, the time-varying reduced-form matrices $\{\mathbf{J}_t, \mathbf{Q}_t, \mathbf{G}_t\}$.[8]

As shown in Cagliarini and Kulish (2013), the existence and uniqueness of the terminal regime, together with a rank condition on the time-varying structural matrices, are necessary and sufficient for the existence and uniqueness of a solution under any finite sequence of structural changes. In our case, existence and uniqueness likewise require a well-defined terminal regime, while sufficiency additionally requires that the matrices $(\mathbf{I} - \Lambda_t \mathbf{Q}_{t+1})$ be of full rank.

The logic behind the backward recursion is straightforward. By assuming that non-stationarity stops after a distant terminal date T, the model becomes stationary from that point onward and the terminal regime can be solved using conventional methods. Drawing a parallel with the Turnpike Theorem, one also assumes a terminal condition far in the future and then works backward to determine the earlier path of the economy.

In our setup, the recursion would, of course, be well-defined when the underlying non-linear model features only a finite sequence of structural changes, in which case there is no infinite-horizon counterpart to approximate. When the solution is instead used as an approximation to an infinite-horizon non-stationary system, the Turnpike conditions become important for the accuracy of the piecewise linear solution: they ensure that the finite–horizon solution remains close to the infinite–horizon path for most of the horizon. Without

[8] The solution is piecewise linear in time, but at any given point in time it is conditionally linear in the state y_t. The occasionally binding constraint solution of Guerrieri and Iacoviello (2015) and Jones (2017) can also be seen as an iterative application of the recursions (16)–(18), where iterations determine the duration of regimes consistent with occasionally binding constraints. They describe their solution as piecewise linear because the reduced forms are tied to the state and constraints, rather than purely to time.

such conditions, the recursions (16)–(18) can still be computed but may yield solutions that fail to approximate the infinite-horizon counterpart. For example, it is well known that the recursion for $\mathbf{Q}_t$ can generate explosive dynamics as the length of a temporary fixed-interest-rate regime increases, even though such a path does not approximate the equilibrium of an economy with a permanently fixed interest rate, a result known as the forward guidance puzzle.

The linear rational expectations framework in (13) is flexible enough to accommodate both anticipated and unanticipated structural changes. To establish the recursions (16)–(18), we used the fact that agents anticipate the evolution of the structural equations, taking into account the expected reduced-form solution at that point in time, namely $\mathbb{E}_t y_{t+1} = \mathbf{J}_{t+1} + \mathbf{Q}_{t+1} y_t$. Accordingly, the recursions link the reduced form matrices at t with $\mathbf{J}_{t+1}$, $\mathbf{Q}_{t+1}$, and $\mathbf{G}_{t+1}$.

If instead the structural changes are assumed to be unanticipated, the solution to (13) is still a time-varying coefficients VAR, but the recursions collapse to depend only on current-period information:

$$\mathbf{J}_t = (\mathbf{I} - \Lambda_t \mathbf{Q}_t)^{-1} (\mathbf{\Gamma}_t + \Lambda_t \mathbf{J}_t), \tag{19}$$

$$\mathbf{Q}_t = (\mathbf{I} - \Lambda_t \mathbf{Q}_t)^{-1} \Phi_t, \tag{20}$$

$$\mathbf{G}_t = (\mathbf{I} - \Lambda_t \mathbf{Q}_t)^{-1} \Psi_t. \tag{21}$$

The choice between anticipated and unanticipated structures is not just technical. For the time-varying coefficient VAR to provide an accurate approximation to the global non-linear solution, the information available to agents in the approximation must match the information available in the underlying non-linear model. Consistency of the information structure is therefore crucial for the accuracy of the piecewise linear solution.

Up to this point, we have shown that when the structural matrices are time-varying, the reduced-form solution takes the form of a time-varying VAR. Information consistency, as we illustrate in the numerical results below, is crucial for ensuring that the piecewise linear solution aligns with the global non-linear solution. Equally important is the choice of expansion points around which the first-order Taylor approximation is taken, since these directly shape the coefficients of the structural matrices and, in turn, the reduced-form time-varying VAR. In the next section, we turn to the role of expansion points in determining the accuracy of the piecewise linear approximation.

4 Expansion Points

An important consideration when linearizing non-linear non-stationary models, where no well-defined steady state exists, is the choice of expansion points – or more precisely, expansion sequences – around which to carry out the linearization. A natural candidate is the path the non-linear economy would have taken in the absence of stochastic shocks. For example, in a structural transformation model, Buera et al. (2023) refer to this path as the *stable transformation path*.[9] In the application below with the non-stationary neoclassical growth model, we refer to this path as the *growth path*: the trajectory the economy would follow if $\varepsilon_{Z,t} = 0$ and $Z_t = 1$ for all t. Formally, it solves

$$(C_t^*)^{-\gamma} = \beta \left[(C_{t+1}^*)^{-\gamma} R_{t+1}^* \right], \tag{22}$$

$$R_t^* = 1 - \delta + \alpha \left(K_{t-1}^*\right)^{\alpha-1} A_t^{1-\alpha}, \tag{23}$$

$$C_t^* + K_t^* = (1-\delta)K_{t-1}^* + \left(K_{t-1}^*\right)^{\alpha} A_t^{1-\alpha}, \tag{24}$$

subject to the terminal condition $K_{T+1}^* = A_{T+1}k^*$, where k^* corresponds to the steady state of the detrended system.

This particular terminal condition, K_{T+1}^*, is not crucial under the conditions of the Turnpike Theorem. When those conditions hold, any sufficiently distant terminal condition has no effect on the approximation for the horizon of interest.

In the literature, some authors have instead used a sequence of steady states as the expansion sequence – an approach we refer to as the 'naive' approach. In our application, this amounts to computing the steady state that would prevail at each value of the non-stationary productivity parameter. For the case of growing productivity, the naive path is given by

$$K_{t-1}^* = \left(\frac{1-(1-\delta)\beta}{\alpha\beta A_t^{1-\alpha}} \right)^{\frac{1}{\alpha-1}}, \tag{25}$$

$$R_t^* = 1 - \delta + \alpha \left(K_{t-1}^*\right)^{\alpha-1} A_t^{1-\alpha}, \tag{26}$$

$$C_t^* = \left(K_{t-1}^*\right)^{\alpha} A_t^{1-\alpha} - \delta K_{t-1}^*. \tag{27}$$

Any expansion sequence that does not coincide with the non-stochastic path the economy would follow in the absence of shocks induces a first-order error in the approximation.

[9] This path is the correct expansion basis for extending structural transformation models into stochastic environments.

As we formally show in the appendix, expansions around alternative sequences – such as the naive path – cannot eliminate this bias. In contrast, using the non-stochastic path as the expansion basis ensures that the linearized system is locally accurate to first order, making it the appropriate benchmark for the piecewise linear solution.

To further illustrate the role of the expansion basis and to introduce notation, for any variable X_t let X_t^* denote its expansion point. We begin with the stationary case in which $A_t = A$ for all t, so that the model in (6)–(8) admits a well-defined steady state in levels, denoted by $\{C^*, K^*, R^*\}$. In the stationary case, the expansion point coincides with the steady state in levels.

Log-linearizing the first-order condition (6) gives

$$\gamma\left(\log C_t - \log C^*\right) = \gamma\left(\mathbb{E}_t \log C_{t+1} - \log C^*\right) - \left(\mathbb{E}_t \log R_{t+1} - \log R^*\right). \tag{28}$$

In much of the literature, it is standard to work with percentage deviations from steady state, $\widehat{X}_t = \log X_t - \log X^*$. In our case, however, it is preferable to work directly with log variables, since a well-defined steady state may not exist and the expansion point may itself be time-varying. Moreover, working in deviations from a time-varying expansion sequence would mask precisely the issue we wish to highlight.

In the non-stationary case where $A_t = gA_{t-1}$ with $g > 1$, it may seem natural to proceed as in the stationary case and linearize at time t around $\{C_t^*, K_t^*, R_t^*\}$. This yields

$$\gamma\left(\log C_t - \log C_t^*\right) = \gamma\left(\mathbb{E}_t \log C_{t+1} - \log C_t^*\right) - \left(\mathbb{E}_t \log R_{t+1} - \log R_t^*\right). \tag{29}$$

Equation (29) may provide a good approximation to the Euler equation in period t. However, notice that the treatment of the approximation point for $\log C_{t+1}$ is inconsistent across periods. To see this, consider linearizing at time $t+1$ around $\{C_{t+1}^*, K_{t+1}^*, R_{t+1}^*\}$:

$$\gamma\left(\log C_{t+1} - \log C_{t+1}^*\right) = \gamma\left(\mathbb{E}_{t+1} \log C_{t+2} - \log C_{t+1}^*\right) - \left(\mathbb{E}_{t+1} \log R_{t+2} - \log R_{t+1}^*\right).$$

The inconsistency is clear: in period t, the Euler equation involves the log-deviation of C_{t+1} from C_t^*, whereas in period $t+1$ it is measured relative to C_{t+1}^*. Since these equations are derived from (6), we conjecture that treating expansion points across time in a consistent manner yields a more accurate approximation to (6)–(8) along a given path.

The same logic applies more generally. In models with adjustment costs, habit formation, or other features that bring more lags or leads of variables into the same equilibrium condition, consistency requires that each dated variable be expanded around its own dated expansion point. For example, if consumption were to appear in an equation at $t-1$, t, and

$t+1$, then the correct first-order approximation must use $\{C^*_{t-1}, C^*_t, C^*_{t+1}\}$. Any cross-dating of expansion points – such as measuring C_{t+1} around C^*_t – introduces a first-order error and undermines the accuracy of the approximation.

Therefore, in our case a time-consistent treatment of expansion points for the Euler equation (6) requires linearizing at time t around $\{C^*_t, C^*_{t+1}, R^*_{t+1}\}$, which yields

$$\gamma\left(\log C_t - \log C^*_t\right) = \gamma\left(\mathbb{E}_t \log C_{t+1} - \log C^*_{t+1}\right) - \left(\mathbb{E}_t \log R_{t+1} - \log R^*_{t+1}\right), \tag{30}$$

where C^*_t, C^*_{t+1}, and R^*_{t+1} denote the corresponding expansion points.

Next, consider the impact of relying on (29) rather than (30) when constructing the structural matrices in (13). For the Euler equation, this choice does not affect the rows of $\mathbf{A}_t$, $\mathbf{B}_t$, $\mathbf{D}_t$, or $\mathbf{F}_t$. It does, however, have a first-order effect on the row of $\mathbf{C}_t$ corresponding to that equation. To see this, notice that with time-inconsistent expansion points, i.e. using (29), we obtain

$$\mathbf{C}_t = \begin{bmatrix} \log R^*_t \\ \cdots \end{bmatrix}, \tag{31}$$

whereas with time-consistent expansion points, i.e. using (30), we arrive at

$$\mathbf{C}_t = \begin{bmatrix} \log C^*_t - \log C^*_{t+1} + \log R^*_{t+1} \\ \cdots \end{bmatrix}. \tag{32}$$

It is clear that these various approaches are not innocuous. While they all ultimately yield the time-varying reduced-form VAR,

$$y_t = \mathbf{J}_t + \mathbf{Q}_t y_{t-1} + \mathbf{G}_t \varepsilon_t,$$

the matrices $\mathbf{J}_t$, $\mathbf{Q}_t$, and $\mathbf{G}_t$ depend crucially on three factors: (i) the choice of expansion sequence, (ii) whether expansion points are applied in a time-consistent way, and (iii) whether agents anticipate structural changes. Taken together, these three factors imply eight possible implementations of the piecewise linear solution. As we show below, the solution in (15) provides an accurate approximation when it is based on the growth path, uses time-consistent expansion points, and assumes that agents anticipate these changes.

In the next section, we evaluate these alternative implementations in the context of the non-stationary stochastic growth model. By comparing their performance against accurate global non-linear benchmarks, we show how the choice of expansion sequence, the treatment of time consistency of the expansion points, and the information structure jointly determine the accuracy of the piecewise linear solution.

5 Numerical results

We now evaluate the accuracy of the alternative implementations of the piecewise linear solution discussed in Section 4. Our benchmark is the global non-linear solution in the case with a stationary representation and the extended function path (EFP) method of Maliar et al. (2020) in the case without one. The comparison allows us to quantify how the choice of expansion path, the time consistency of approximation points, and the information structure affect the accuracy of the time-varying VAR solution. In addition, we assess the sensitivity of the best-performing implementation to parameters that govern the degree of non-linearity in the model, such as risk aversion, the growth rate of productivity, and the variance of shocks.[10]

5.1 The case of $Z_t K_{t-1}^{\alpha} A_t^{1-\alpha}$ with $A_t = gA_{t-1}$

In our first experiment we fix parameters at $\alpha = 0.36$, $\beta = 0.99$, $\gamma = 1$, $\delta = 0.025$, $\rho = 0.95$, $\sigma = 0.01$, $A_0 = 1$, $g = 1.01$, and $T = 450$. We solve the non-stationary model using alternative sequences of structural matrices. Exploiting the property of balanced growth, we then transform the model into a stationary representation and compute the global non-linear solution using a projection algorithm following Krueger and Kubler (2004) and Judd et al. (2014). After obtaining the accurate stationary solution, we recover the implied non-stationary solution and compare it with the piecewise linear solutions in (15) under different expansion sequences and information assumptions. For comparability, we use the same initial conditions for capital and productivity, as well as the same sequence of productivity shocks, across all simulations.

Figure 2 compares the infinite-horizon global non-linear solution with the most accurate implementation of the piecewise linear method. The top panel plots the level of capital, K_t, while the bottom panel shows detrended capital, $k_t = K_t/A_t$. For most of the horizon, the two solutions are nearly indistinguishable, confirming that the piecewise linear method closely tracks the global benchmark. Toward the end of the simulation, small differences emerge because the terminal conditions differ: in the detrended infinite-horizon model there is no well-defined terminal level of A_t, whereas the piecewise linear solution requires assuming a stationary terminal level. The gap becomes more pronounced in the final 50 periods as the piecewise linear economy approaches its terminal regime. As noted

[10] All results are obtained using Matlab 2023a on a laptop with an Apple M2 chip (3.5GHz) and 8GB of RAM.

Figure 2: Piecewise linear versus global non-linear solution.

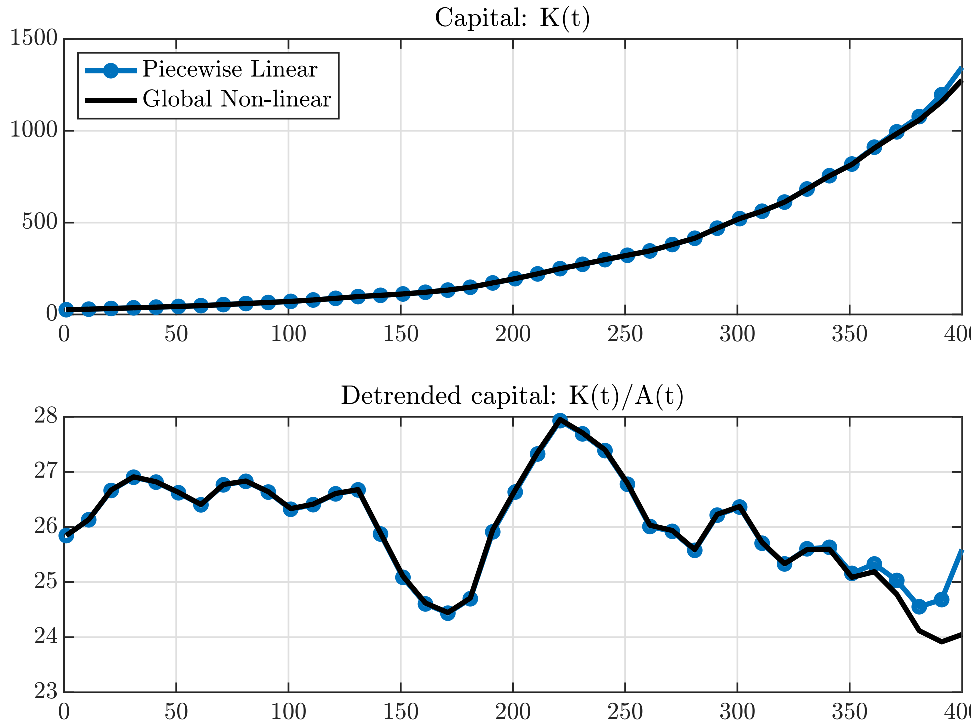

Notes: The piecewise linear solution is constructed using the anticipated growth path and time-consistent expansion points.

by Maliar et al. (2020), such tail inaccuracies typically affect the last 100–150 periods without undermining the quality of the approximation over the relevant range.

Table 1 reports running times as well as mean and maximum absolute differences (in percent) between the global solution and each piecewise linear implementation, averaged across 100 simulations. To assess the impact of the terminal regime, we compute errors over different horizons: the first 50, 100, 250, 400, and 450 periods. As is common in the computational literature, we measure accuracy using the unit-free absolute difference

$$\epsilon_t = \log_{10}\left|\frac{K_t^{\text{PW}}}{K_t^{\text{Global}}} - 1\right|,$$

where K_t^{PW} and K_t^{Global} denote capital under the piecewise linear and global solutions, respectively. But for ease of interpretation, in the tables we report the corresponding percentage deviations rather than the $\log_{10}$ values themselves. For example, $\epsilon = -2.0$ corresponds to a mean error of about $10^{-2.0} \approx 1.0\%$, while $\epsilon = -3.5$ corresponds to 0.03%.

As the last column of Table 1 shows, the piecewise linear solution that relies on the growth path as the expansion sequence, assumes that agents anticipate productivity changes (that is, the sequence of time-varying structural parameters in the linearized equations induced by growing productivity), and uses time-consistent expansion points is highly accurate, with mean errors below one-tenth of a percent.

The table also quantifies how departing from these choices affects accuracy. In some cases, mean errors in the first 350 periods can be as large as $10^{-0.163} \approx 68\%$, as occurs when using the naive path with unanticipated changes and two approximation points. Comparing the last two columns highlights the importance of time consistency: using the anticipated growth path but only one expansion point yields a worse approximation than the unanticipated naive path with a single expansion point.

This illustrates that deviating from the correct implementation of the piecewise linear solution can introduce errors that sometimes offset each other. For instance, while assuming unanticipated changes is inconsistent with the information structure of the non-linear model – where agents are assumed to anticipate the path of $A_t = gA_{t-1}$ – it can nevertheless deliver a smaller error than some other inconsistent cases. This, of course, is not a general result but depends on the particular application.

Figure 3 illustrates the implications of different implementations of the piecewise linear solution under the growth path assumption, showing one simulation for detrended capital, $k_t = K_t/A_t$. The four panels correspond to cases with anticipated versus unanticipated changes and with one versus two approximation points. The figure highlights how inconsistent treatments of expansion points or information induce a first-order error in levels: some implementations systematically increase k_t, while others decrease it. This is consistent with the theoretical discussion above, where we showed that using time-inconsistent or misspecified expansion sequences shifts the structural matrices in a way that biases the reduced-form solution.

But the figure also suggests that these errors may be less severe when comparing first differences of log capital, $\log(K_t/K_{t-1})$, rather than levels, as all solutions seem to track the ups and downs of the global solution. Differencing removes the first-order error in levels, leaving behind a smaller residual discrepancy. In the next table we therefore report results based on $\log(K_t/K_{t-1})$, which provide a complementary measure of the accuracy of the piecewise linear solution.

This observation is particularly relevant for empirical applications. When estimating non-stationary models, the choice of observables matters: if estimation relies on level vari-

Table 1: Comparison of piecewise linear solutions for K_t: the case of A_t.

Expansion Sequences	Naive path				Growth path			
Information	Unanticipated		Anticipated		Unanticipated		Anticipated	
Expansion Points	1 ap.p.	2 ap.p.	1 ap.p.	2 ap.p.	1 ap.p.	2 ap.p.	1 ap.p.	2 ap.p.
Mean errors across t periods in percents								
$t \in [0, 50]$	5.03	33.65	2.35	21.09	11.0	6.96	15.4	0.02
$t \in [0, 100]$	7.28	49.56	3.32	30.44	14.9	9.51	20.7	0.03
$t \in [0, 350]$	9.99	68.72	4.53	41.45	19.1	12.2	26.3	0.04
$t \in [0, 400]$	10.16	69.87	4.57	42.19	19.4	12.4	26.6	0.04
$t \in [0, 450]$	10.28	70.76	4.19	43.27	19.6	12.5	26.8	0.06
Maximum errors across t periods in percents								
$t \in [0, 50]$	10.33	66.12	5.30	40.57	18.9	12.0	25.9	0.14
$t \in [0, 100]$	12.35	78.10	6.40	47.25	21.4	13.4	29.0	0.23
$t \in [0, 350]$	13.19	80.87	6.89	48.99	21.9	13.6	29.7	0.34
$t \in [0, 400]$	13.19	80.87	6.89	49.79	21.9	13.6	29.7	0.34
$t \in [0, 450]$	13.22	80.87	6.89	62.92	21.9	13.6	29.7	5.72
Running time, in seconds								
Solution	0.21	0.22	0.22	0.25	0.18	0.21	0.22	0.23
Simulation	0.15	0.12	0.12	0.12	0.12	0.12	0.13	0.12
Total	0.44	0.36	0.34	0.34	0.37	0.30	0.34	0.35

Notes: "Mean errors" and "Maximum errors" are, respectively, mean and maximum unit-free absolute difference between the exact solution for capital and the solution delivered by a method in the column. The difference between the solutions is computed across 100 simulations.

Figure 3: Various piecewise linear solutions under the growth path assumption.

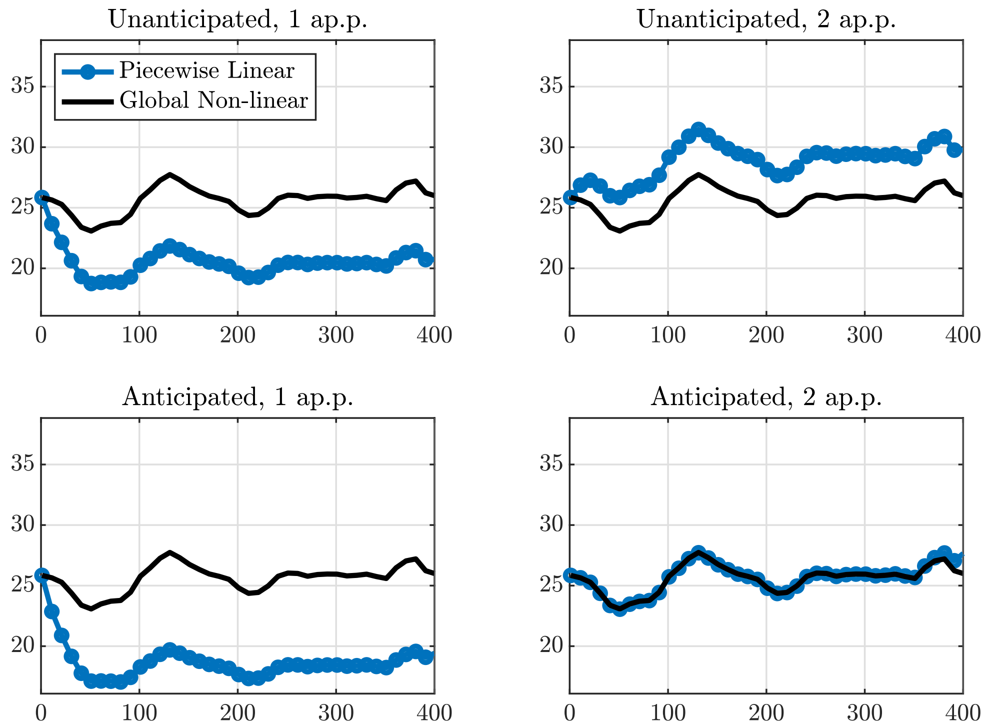

ables while using an inaccurate implementation of the piecewise linear solution, the induced first-order errors can contaminate the mapping between theory and data. But if observables are in growth rates or first differences, these errors may be smaller as we show next, reducing potential biases. This may explain why it may still be possible to obtain reasonable parameter estimates despite using an inaccurate underlying solution. The best option, however, is to rely on the accurate implementation of the piecewise linear method, which avoids this issue altogether.

In Table 2 we report mean and maximum errors for the growth rate of capital. Again, it is clear that the solution constructed using two approximation points from the anticipated growth path is very accurate: its mean errors range from 0.08% to 0.26%. The remaining piecewise linear solutions share a similar pattern: their inaccuracy is concentrated at the beginning of the simulation and becomes much smaller toward the end, and is smaller than for the level of capital in all cases. Figure 3 helps explain why the error is more pronounced at the beginning. The discrepancy arises because the capital level implied by the accurate solution differs from that implied by inaccurate time-varying VARs. Since the simulation

Table 2: Comparison of piecewise linear solutions for $\log(K_t/K_{t-1})$.

	Naive path				Growth path			
	Unanticipated		Anticipated		Unanticipated		Anticipated	
	1 ap.p.	2 ap.p.	1 ap.p.	2 ap.p.	1 ap.p.	2 ap.p.	1 ap.p.	2 ap.p.
Mean errors across t periods in percents								
$t \in [0, 50]$	15.39	87.19	6.71	57.87	35.72	18.92	51.41	0.08
$t \in [0, 100]$	6.93	31.21	3.93	20.37	12.59	5.98	18.42	0.11
$t \in [0, 350]$	3.01	3.88	2.55	3.39	1.49	0.49	1.70	0.14
$t \in [0, 400]$	2.87	3.45	2.52	3.21	1.33	0.42	1.52	0.17
$t \in [0, 450]$	2.78	3.17	2.90	3.76	1.22	0.38	1.74	0.26
Maximum errors across t periods in percents								
$t \in [0, 50]$	104.29	439.42	18.11	297.57	174.49	96.26	251.80	9.11
$t \in [0, 100]$	104.29	439.42	39.26	297.57	174.49	96.26	251.80	10.29
$t \in [0, 350]$	104.29	439.42	39.26	297.57	174.49	96.26	251.80	12.91
$t \in [0, 400]$	104.29	439.42	39.26	297.57	174.49	96.26	251.80	12.91
$t \in [0, 450]$	104.29	439.42	89.35	297.57	174.49	96.26	251.80	61.65
Running time, in seconds								
Solution	0.21	0.22	0.22	0.25	0.18	0.21	0.22	0.23
Simulation	0.15	0.12	0.12	0.12	0.12	0.12	0.13	0.12
Total	0.44	0.36	0.34	0.34	0.37	0.30	0.34	0.35

Notes: "Mean errors" and "Maximum errors" are, respectively, mean and maximum unit-free absolute difference between the exact solution for the growth rate of capital and the solution delivered by a method in the column. The difference between the solutions is computed across 100 simulations.

starts from the accurate capital level, it takes time to adjust to the implied one. During this adjustment, the difference in capital growth rates is large, as seen in the table. Afterward, the piecewise methods track fluctuations in the growth rate of capital reasonably well.[11]

Constructing piecewise linear solutions is computationally inexpensive: in our experiments, solving the model and simulating the solution 100 times took less than a second. About 70% of the total running time is spent on solving for the decision rules and the remaining 30% on simulating paths. Because the method constructs a path of decision functions that are piecewise linear in time, the model can be simulated many times at very low cost, which makes this approach attractive for estimation.[12] In our application, computing the growth path was straightforward and equally cheap. In more complex models, however, constructing the non-stochastic path may require solving a high-dimensional non-linear system, which can be computationally intensive depending on the setting. This step can be carried out with standard solvers for deterministic systems of non-linear equations.[13] Once the non-stochastic path is available, by contrast, obtaining the time-varying VAR is not computationally burdensome. Moreover, in richer non-stationary models with many state variables, global non-linear solution methods are typically infeasible. In such cases, even if computing the non-stochastic path is computationally demanding, the time-varying VAR approach provides a feasible way forward, and this paper shows how to implement it.

The accuracy of all time-varying VARs deteriorates toward the end of the simulation. This occurs because we approximate the solution to an infinite-horizon economy with a finite-horizon counterpart, or equivalently, with an economy in which productivity grows but eventually stalls. This type of approximation has been extensively studied in the *turnpike* literature: see Dorfman et al. (1958), Brock (1971), McKenzie (1976), Majumdar and Zilcha (1987), and Mitra and Nyarko (1991) for theoretical results, and Maliar et al. (2020) for some recent numerical results. The main conclusion of this research is that a solution to the finite-horizon model can approximate the infinite-horizon solution arbitrarily well, provided the horizon is sufficiently long. In our context, this means that as T increases, the inaccuracy observed in the first 450 periods of the non-stationary model disappears.

[11]From the perspective of empirical work, this feature has an important implication. In estimation, the data can be thought to 'discipline' the solution: by adjusting initial conditions, observed growth rates can be matched even when the underlying piecewise-linear solution is inaccurate in levels.

[12]Kulish and Pagan (2017) show how to set up the full-information estimation of models with structural changes.

[13]These include solvers based based on Newton style, Gauss-Siedel or Gauss-Jacobi iterative methods

In the next exercise we assess how the accuracy of the correct implementation of the piecewise linear solution varies with parameters that govern the degree of non-linearity in the model: the curvature of the utility function, γ, the growth rate of productivity, g, and the volatility of the stochastic productivity shock, σ. Specifically, we consider $\gamma \in \{0.1, 1, 3\}$, $\sigma \in \{0.01, 0.03\}$, and $g \in \{0.01, 0.03\}$. Across these parametrizations, the accuracy of the piecewise linear method with two approximation points along the growth path is similar to that reported in Table 1, and the total running time remains under one second in all cases. As expected, accuracy deteriorates when the variance of shocks increases, reflecting the certainty equivalence property of the piecewise linear solution, and when the coefficient of relative risk aversion, γ, rises. Although the piecewise linear solution is slightly less accurate in these cases, it still provides a very close approximation to the global solution, with errors that do not exceed half a percent.

5.2 The case of $Z_t K_t^{\alpha_t}$ with $\alpha_t = \alpha_0 + m_\alpha t$

The case with growing labor-augmenting productivity was useful because it allowed us to compare the accuracy of the piecewise linear solution for the non-stationary model in levels with the global non-linear solution of its stationary representation – effectively benchmarking against the computational gold standard for solving non-linear dynamic stochastic rational expectations models. In practice, however, one would not rely on the piecewise linear solution if a stationary representation were available.

We therefore turn to a case that does not admit such a representation: a drifting capital share, α_t. Specifically, we consider a version of model (1)–(3) in which the parameter α_t grows for the first 500 periods from an initial value of 0.3 up to 0.5. Because this model lacks a stationary representation, we benchmark the accuracy of the piecewise linear solution against the Extended Function Path (EFP) solution proposed by Maliar et al. (2020).

An additional complication in this case is that there is no closed-form expression for the growth path. This is also the more realistic situation in practice, since models with arbitrary forms of non-stationarity typically do not admit simple analytical representations of their non-stochastic paths. To obtain a reliable piecewise linear solution, we therefore solve the following system of non-linear equations:

$$
\begin{aligned}
(C_t^*)^{-\gamma} &= \beta \left[(C_{t+1}^*)^{-\gamma} R_{t+1}^* \right], \\
R_t^* &= 1 - \delta + \alpha_t \, (K_{t-1}^*)^{\alpha_t - 1} A_t^{1-\alpha_t}, \\
C_t^* + K_t^* &= (1-\delta) K_{t-1}^* + (K_{t-1}^*)^{\alpha_t} A_t^{1-\alpha_t}.
\end{aligned}
$$

Table 3: Sensitivity analysis for best piecewise linear solution: A_t case.

Parameters	Calib 1	Calib 2	Calib 3	Calib 4	Calib 5	Calib 6
γ	1	1	1	1	0.1	3
σ_ϵ	0.01	0.01	0.03	0.01	0.01	0.01
g	1.00	1.01	1.01	1.03	1.01	1.01
Mean errors across t periods in percents						
$t \in [0, 50]$	0.01	0.02	0.15	0.02	0.02	0.13
$t \in [0, 100]$	0.02	0.03	0.25	0.03	0.02	0.19
$t \in [0, 350]$	0.05	0.04	0.39	0.04	0.03	0.29
$t \in [0, 400]$	0.05	0.04	0.38	0.04	0.03	0.27
$t \in [0, 450]$	0.05	0.06	0.41	0.06	0.03	0.35
Maximum errors across t periods in percents						
$t \in [0, 50]$	0.13	0.14	1.21	0.15	0.08	0.40
$t \in [0, 100]$	0.22	0.23	2.04	0.22	0.11	0.61
$t \in [0, 350]$	0.35	0.34	3.07	0.31	0.11	0.86
$t \in [0, 400]$	0.35	0.34	3.07	0.31	0.11	0.86
$t \in [0, 450]$	0.35	5.72	5.60	12.66	2.10	14.44
Running time, in seconds						
Solution	0.25	0.26	0.24	0.25	0.24	0.24
Simulation	0.12	0.13	0.12	0.13	0.12	0.12
Total	0.37	0.39	0.37	0.37	0.36	0.36

Notes: "Mean errors" and "Maximum errors" are, respectively, mean and maximum unit-free absolute difference between the exact solution for capital and the piecewise linear solution delivered under the parameterization in the column. Each column corresponds to a different calibration of the model (Calib 1 to Calib 6), with parameter values shown in the top rows. The difference between the solutions is computed across 100 simulations. The time horizon is $T = 450$, and the terminal condition is constructed by using the T-period stationary economy in all experiments.

We assume $\alpha_0 = 0.3$ and $\alpha_T = 0.5$ for the initial and terminal conditions of the capital share. Our Matlab implementation solves the system in under one second.[14]

Figure 4 plots the capital stock, K_t, under the piecewise linear solution alongside the accurate non-linear solution obtained with the EFP method. The two paths are nearly indistinguishable over most of the horizon. Small differences emerge only toward the end, reflecting the fact that the terminal conditions differ: in the EFP solution the terminal condition is implied by the non-linear model at $\alpha = 0.5$, whereas in the piecewise linear solution it is given by the standard linear rational expectations solution at $\alpha = 0.5$.

Figure 4: $\alpha = [0.3, 0.5]$

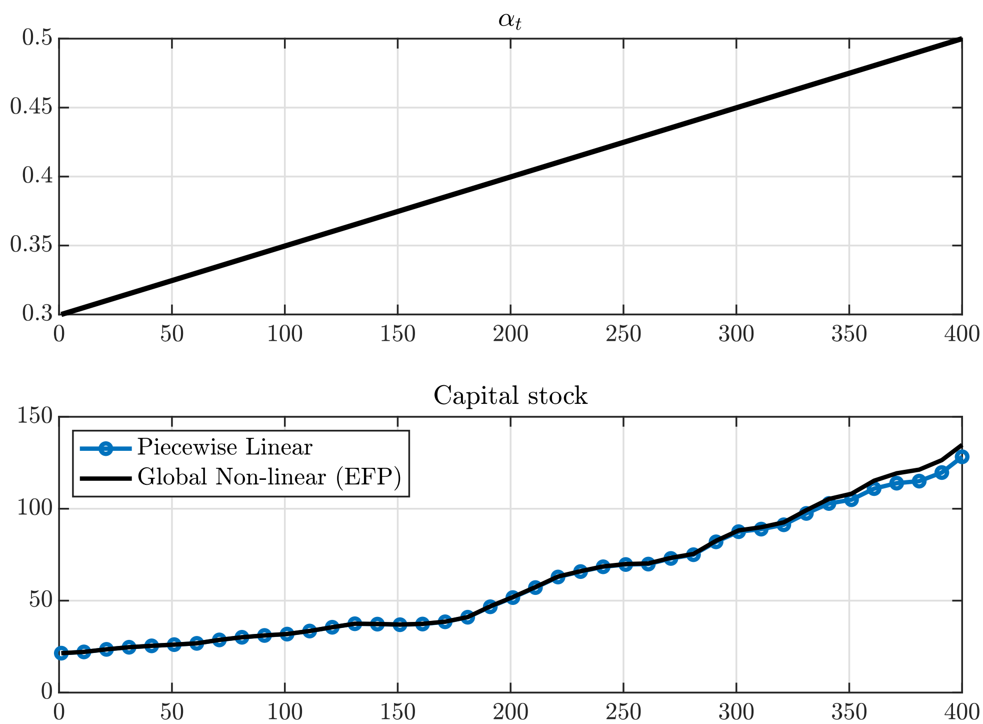

Table 4 reports mean and maximum errors for the model with a drifting capital share, α_t. As in the case with growing productivity, the most accurate implementation is the one that combines the growth path, time-consistent expansion points, and anticipated changes, with mean errors consistently below 0.1%. However, compared to the A_t case, the other implementations perform much better: mean errors in the naive path with anticipated changes are

[14]We use standard solvers for non-linear systems, such as Newton-type methods or Gauss-Seidel/Gauss-Jacobi iterations.

only around 2–4% for the first 350 periods, and even some unanticipated cases stay within single-digit errors. This suggests that, in settings where computing the growth path is costly, a simpler approximation such as the naive path may still yield reasonable accuracy. From a numerical standpoint, this is an encouraging result: it indicates that the structure of non-stationarity matters, and that it may be possible to exploit such differences when choosing approximation strategies.

Table 4: Comparison of piecewise linear solutions for K_t: the case of α_t.

	Naive path				Growth path			
	Unanticipated		Anticipated		Unanticipated		Anticipated	
	1 ap.p.	2 ap.p.	1 ap.p.	2 ap.p.	1 ap.p.	2 ap.p.	1 ap.p.	2 ap.p.
Mean errors across t periods in percents								
$t \in [0, 50]$	0.44	4.35	0.02	4.02	1.87	0.39	1.14	0.02
$t \in [0, 100]$	0.62	6.36	0.07	6.14	3.80	0.48	2.78	0.03
$t \in [0, 350]$	0.67	12.01	0.30	11.33	8.69	0.24	7.38	0.06
$t \in [0, 400]$	0.65	13.15	0.35	12.31	9.37	0.26	7.55	0.08
$t \in [0, 450]$	0.67	14.36	0.40	13.38	9.90	0.28	6.45	0.12
Maximum errors across t periods in percents								
$t \in [0, 50]$	1.11	8.05	0.47	8.25	6.05	0.75	5.09	0.13
$t \in [0, 100]$	1.50	10.52	0.85	10.54	9.36	0.78	8.33	0.22
$t \in [0, 350]$	1.64	23.40	1.66	20.88	16.05	0.78	13.18	0.64
$t \in [0, 400]$	1.64	27.29	1.81	24.26	16.76	0.78	13.18	2.12
$t \in [0, 450]$	2.36	31.48	2.70	29.32	16.79	0.78	13.18	4.49
Running time, in seconds								
Solution	0.67	0.66	0.66	0.67	0.62	0.66	0.66	0.68
Simulation	0.12	0.12	0.12	0.12	0.12	0.12	0.12	0.12
Total	0.79	0.78	0.78	0.79	0.74	0.78	0.78	0.80

Notes: "Mean errors" and "Maximum errors" are, respectively, mean and maximum unit-free absolute difference between the exact solution for capital and the solution delivered by a method in the column. The difference between the solutions is computed across 100 simulations.

Having established the relative performance of alternative implementations in Table 4, Table 5 turns to a sensitivity analysis of the best-performing case – the anticipated growth path with time-consistent expansion points. We vary parameters that govern nonlinearity: risk aversion γ, the volatility of productivity shocks σ, and the terminal value of the capital share α_T. As expected, accuracy declines somewhat when shocks are more volatile ($\sigma =$

0.03) or when risk aversion increases ($\gamma = 3$), with mean errors rising modestly but still remaining below 1%. The piecewise linear solution also deteriorates when the terminal capital share grows to $\alpha_T = 0.6$, again reflecting stronger nonlinearities in the system. Even so, the piecewise linear solution continues to provide a highly accurate and computationally feasible approximation, with running times under a second. These results reinforce the conclusion that the method is reliable across a broad range of parameterizations, while also highlighting how the degree of nonlinearity shapes the limits of its accuracy.

Table 5: Sensitivity analysis for best piecewise linear solution: α_t case.

Parameters	Calib 1	Calib 2	Calib 3	Calib 4	Calib 5	Calib 6
γ	1	1	1	1	0.3	3
σ_ϵ	0.01	0.01	0.03	0.01	0.01	0.01
α_T	0.3	0.5	0.5	0.6	0.5	0.5
Mean errors across t periods in percents						
$t \in [0, 50]$	0.02	0.02	0.14	0.02	0.01	0.10
$t \in [0, 100]$	0.03	0.03	0.24	0.03	0.02	0.16
$t \in [0, 350]$	0.05	0.06	0.42	0.16	0.03	0.42
$t \in [0, 400]$	0.05	0.08	0.50	0.28	0.03	0.62
$t \in [0, 450]$	0.05	0.12	0.62	0.47	0.05	0.90
Maximum errors across t periods in percents						
$t \in [0, 50]$	0.14	0.13	1.17	0.13	0.09	0.34
$t \in [0, 100]$	0.22	0.22	1.99	0.22	0.17	0.54
$t \in [0, 350]$	0.37	0.64	3.26	9.73	0.20	5.63
$t \in [0, 400]$	0.37	2.12	3.59	27.50	0.30	16.69
$t \in [0, 450]$	0.37	4.49	6.89	33.32	1.49	21.45
Running time, in seconds						
Solution	0.44	0.67	0.68	0.75	0.81	0.96
Simulation	0.12	0.12	0.12	0.12	0.12	0.12
Total	0.56	0.78	0.79	0.87	0.93	1.08

Notes: "Mean errors" and "Maximum errors" are, respectively, mean and maximum unit-free absolute difference between the exact solution for capital and the correct piecewise solution under the parameterization in the column. Each column corresponds to a different calibration of the model (Calib 1 to Calib 6), with parameter values shown in the top rows.

Taken together, the numerical exercises demonstrate that the piecewise linear solution is both accurate and computationally efficient when implemented correctly. Across different sources of non-stationarity, and under a range of parameterizations, the method delivers results that closely track global non-linear benchmarks.

6 Conclusion

We have studied the accuracy and efficiency of the piecewise linear solution of Kulish and Pagan (2017) applied to non-stationary rational expectations models. By benchmarking against global non-linear solutions we have shown that the piecewise linear approach can deliver highly accurate approximations when implemented correctly.

The key requirement is to construct expansion sequences along the non-stochastic path, to apply expansion points in a time-consistent manner, and to model expectations in a way that is consistent with the information structure of the underlying non-linear economy. When these conditions are satisfied, the piecewise linear solution tracks global benchmarks with mean errors below a tenth of a percent. At the same time, the method is computationally efficient, making it particularly attractive for estimation and policy analysis of medium- and large-scale non-stationary models where global methods are infeasible.

Our numerical results also show how deviations from the correct implementation – using naive expansion paths, time-inconsistent approximation points, or unanticipated structural changes – give rise to systematic errors. These findings highlight that implementation choices are not innocuous: accuracy depends critically on the expansion basis and the underlying information assumptions. Finally, sensitivity analysis confirms that while accuracy deteriorates as the degree of nonlinearity increases, the method continues to provide a reliable and feasible approximation across a broad range of parameterizations.

Taken together, the results establish the piecewise linear method as a practical computational strategy for solving non-stationary models with rational expectations. Beyond the examples considered here, the approach opens the door to applications in richer settings without balanced growth paths, including stochastic models of structural transformation and other forms of persistent non-stationarity.

In this paper we have considered exogenous deterministic forms of non-stationarity, such as $A_t = gA_{t-1}$ and $\alpha_t = \alpha_0 + m_\alpha t$. A natural extension could consider exogenous stochastic trends, for example a productivity process following a unit root with drift. In such cases, the non-stochastic path would need to be recomputed as shocks unfold, making the method computationally more demanding but possibly still providing a feasible way forward. Exploring this extension is a promising direction that we leave for future research.

References

Ajevskis, Viktors, "Semi-global solutions to DSGE models: Perturbation around a deterministic path," *Studies in Nonlinear Dynamics & Econometrics*, 2017, *21* (2), 20160065.

Anderson, Gary, "A Reliable and Computationally Efficient Algorithm for Imposing the Saddle Point Property in Dynamic Models," , unpublished manuscript, Board of Governors of the Federal Reserve System 1997.

Binder, Michael and M Hashem Pesaran, "Multivariate Rational Expectations Models and Macroeconometric Modelling: A Review and Some New Results," in in M Hashem Pesaran and M Wickens, eds., *Handbook of Applied Econometrics: Macroeconomics*, Basil Blackwell Oxford 1995, pp. 139–187.

Blanchard, Olivier Jean and Charles M. Kahn, "The Solution of Linear Difference Models Under Rational Expectations," *Econometrica*, July 1980, *48* (5), 1305–1311.

Brock, William A, "Sensitivity of optimal growth paths with respect to a change in target stocks," in "Contributions to the von Neumann growth model," Springer, 1971, pp. 73–89.

Buera, Francisco, Joseph Kaboski, Marti Mestieri, and Daniel O'Connor, "The Stable Transformation Path," Technical Report, CEPR Discussion Papers 2023.

Cagliarini, Adam and Mariano Kulish, "Solving linear rational expectations models with predictable structural changes," *Review of Economics and Statistics*, 2013, *95* (1), 328–336.

Dorfman, Robert, Paul A Samuelson, and Robert M Solow, *Linear programming and economic analysis*, McGraw-Hill, 1958.

Evans, Rick and Kerk Phillips, "Linearization about the current state: A computational method for approximating nonlinear policy functions during simulation," Technical Report, Brigham Young University, Department of Economics, BYU Macroeconomics and 2015.

Gomez-Gonzalez, Patricia and Daniel M Rees, "Same Spain, less pain?," *European Economic Review*, 2018, *110*, 78–107.

Guerrieri, Luca and Matteo Iacoviello, "OccBin: A toolkit for solving dynamic models with occasionally binding constraints easily," *Journal of Monetary Economics*, 2015, *70*, 22–38.

Jones, Callum, "Unanticipated shocks and forward guidance at the zero lower bound," in "Technical Report," NYU Working Paper, 2017.

Judd, Kenneth L, Lilia Maliar, Serguei Maliar, and Rafael Valero, "Smolyak method for solving dynamic economic models: Lagrange interpolation, anisotropic grid and adaptive domain," *Journal of Economic Dynamics and Control*, 2014, *44*, 92–123.

King, Robert and Mark Watson, "The Solution of Singular Linear Difference Systems Under Rational Expectations," *International Economic Review*, 1998, *39* (4), 1015–1026.

King, Robert G, Charles I Plosser, and Sergio T Rebelo, "Production, growth and business cycles: I. The basic neoclassical model," *Journal of monetary Economics*, 1988, *21* (2-3), 195–232.

Klein, Paul, "Using the Generalized Schur Form to Solve a Multivariate Linear Rational Expectations Model," *Journal of Economic Dynamics and Control*, 2000, *24* (10), 1305–1311.

Krueger, Dirk and Felix Kubler, "Computing equilibrium in OLG models with stochastic production," *Journal of Economic Dynamics and Control*, 2004, *28* (7), 1411–1436.

Kulish, Mariano and Adrian Pagan, "Estimation and Solution of Models with Expectations and Structural Changes," *Journal of Applied Econometrics*, 2017, *32* (2), 255–274.

Majumdar, Mukul and Itzhak Zilcha, "Optimal growth in a stochastic environment: some sensitivity and turnpike results," *Journal of Economic Theory*, 1987, *43* (1), 116–133.

Maliar, Lilia, Serguei Maliar, John B Taylor, and Inna Tsener, "A tractable framework for analyzing a class of nonstationary Markov models," *Quantitative Economics*, 2020, *11* (4), 1289–1323.

McKenzie, Lionel W, "Turnpike theory," *Econometrica: Journal of the Econometric Society*, 1976, pp. 841–865.

Mennuni, Alessandro, Juan F Rubio-Ramírez, and Serhiy Stepanchuk, "Dynamic perturbation," *Review of Economic Studies*, 2025, *92* (2), 1157–1192.

Mitra, Tapan and Yaw Nyarko, "On the existence of optimal processes in non-stationary environments," *Journal of Economics*, 1991, *53* (3), 245–270.

Phillips, Kerk L, "Solving and simulating unbalanced growth models using linearization about the current state," *Economics Letters*, 2017, *151*, 35–38.

Sims, Christopher A., "Solving Linear Rational Expectations Models," *Computational Economics*, 2001, *20*, 1–20.

Uhlig, Harald, "A toolkit for analysing nonlinear stochastic models easily," *Computational methods for the study of dynamic economics, Marimon and Scott eds*, 1999.

Appendix A: Why the Consistent Expansion Path is the Right Linearization Base

Setup. Let $s = \{s_t\}_{t=0}^{\infty}$, $\kappa = \{\kappa_t\}_{t=0}^{\infty}$ and $\varepsilon = \{\varepsilon_t\}_{t=0}^{\infty}$, with $s_t \in \mathbb{R}^n$ (an $n \times 1$ vector of state and jump variables), $\kappa_t \in \mathbb{R}^p$ collects structural parameters or exogenous variables, some of which may be time-varying for a finite period of time, and $\varepsilon_t \in \mathbb{R}^q$ (a $q \times 1$ vector of shocks).

Define

$$\mathcal{S} := \ell^{\infty}(\mathbb{R}^n) \quad = \Big\{ s = (s_t)_{t\geq 0} : s_t \in \mathbb{R}^n,\ \|s\|_\infty := \sup_{t\geq 0} \|s_t\| < \infty \Big\}. \tag{33}$$

$$\mathcal{K} := \ell^{\infty}(\mathbb{R}^p) \quad = \Big\{ \kappa = (\kappa_t)_{t\geq 0} : \kappa_t \in \mathbb{R}^p,\ \|\kappa\|_\infty := \sup_{t\geq 0} \|\kappa_t\| < \infty \Big\}. \tag{34}$$

$$\mathcal{E} := \ell^{\infty}(\mathbb{R}^q) \quad = \Big\{ \varepsilon = (\varepsilon_t)_{t\geq 0} : \varepsilon_t \in \mathbb{R}^q,\ \|\varepsilon\|_\infty := \sup_{t\geq 0} \|\varepsilon_t\| < \infty \Big\}. \tag{35}$$

Equipped with the sup norms $\|\cdot\|_\infty$, $\mathcal{S}$, $\mathcal{K}$ and $\mathcal{E}$ are Banach spaces.

The equilibrium conditions of the model at a given period of time are given by

$$E_t f\big(s_{t+1}, s_t, s_{t-1}, \ldots;\ \kappa_t;\ \varepsilon_t\big) = 0,$$

where "…" denotes any finite set of leads/lags; at $t = 0$ required predetermined variables are fixed by initial conditions. We stack the equilibrium conditions for $t = 0, 1, 2, \ldots$ into $F : \mathcal{S} \times \mathcal{K} \times \mathcal{E} \to \mathcal{S}$ which is an equilibrium operator satisfying

$$F(s, \kappa, \varepsilon) = 0. \tag{36}$$

Consistent expansion path. A consistent expansion path is a deterministic sequence $\bar{s} \in \mathcal{S}$ such that

$$F(\bar{s}, \kappa, 0) = 0 \tag{37}$$

and the linear operator $D_s F(\bar{s}, \kappa, 0) : \mathcal{S} \to \mathcal{S}$ is boundedly invertible.[15]

[15] This regularity ensures local uniqueness of $\bar{s}$ for $\varepsilon = 0$ and supports a first-order perturbation analysis via the Implicit Function Theorem. In applications, F stacks the time-t conditions $f_t(s_t, s_{t-1}, \kappa, \varepsilon_t)$ and $D_s F$ becomes a block bi-diagonal operator composed of Jacobians evaluated along $\bar{s}$.

First-order (piecewise) linearization along $\bar{s}$. Consider small shocks ε and seek $s(\varepsilon)$ solving (36) with $s(0) = \bar{s}$. Taylor's theorem in Banach spaces gives, for s near $\bar{s}$ and small ε,

$$F(s,\kappa,\varepsilon) = F(\bar{s},\kappa,0) + D_sF(\bar{s},\kappa,0)[\,s-\bar{s}\,] + D_\varepsilon F(\bar{s},\kappa,0)[\,\varepsilon\,] + R_2(s-\bar{s},\varepsilon), \tag{38}$$

where $\|R_2(\Delta s,\varepsilon)\| = \mathcal{O}\big(\|\Delta s\|^2 + \|\varepsilon\|^2 + \|\Delta s\|\,\|\varepsilon\|\big)$ and $\Delta s \equiv s - \bar{s}$. Imposing (36) and using (37) to remove the constant term yields the variational (first-order) system

$$D_sF(\bar{s},\kappa,0)\big[\,\Delta s\,\big] \;+\; D_\varepsilon F(\bar{s},\kappa,0)\big[\,\varepsilon\,\big] \;=\; 0, \tag{39}$$

whose unique solution is

$$\Delta s^\star(\varepsilon) \;=\; -\big(D_sF(\bar{s},\kappa,0)\big)^{-1} D_\varepsilon F(\bar{s},\kappa,0)[\,\varepsilon\,]. \tag{40}$$

Define the piecewise linear approximation $s^{\mathrm{PL}}(\varepsilon) \equiv \bar{s} + \Delta s^\star(\varepsilon)$, i.e., the model is linearized along the deterministic path $\bar{s}$ so that Jacobians may vary with t. Substituting $s^{\mathrm{PL}}(\varepsilon)$ into (38) gives the residual

$$F\Big(s^{\mathrm{PL}}(\varepsilon),\kappa,\varepsilon\Big) \;=\; R_2(\Delta s^\star(\varepsilon),\varepsilon) \;=\; \mathcal{O}\big(\|\varepsilon\|^2\big)\,, \tag{41}$$

so the piecewise linear solution is first-order accurate. By the Implicit Function Theorem, the true solution admits the expansion $s(\varepsilon) = \bar{s} + \Delta s^\star(\varepsilon) + o(\|\varepsilon\|)$, validating the construction.

Why other expansion sequences are wrong at first order. Let $\tilde{s}$ be any sequence with $F(\tilde{s},\kappa,0) \neq 0$. Repeating (38) around $\tilde{s}$ (with ε small and s near $\tilde{s}$) gives

$$F(s,\kappa,\varepsilon) = \underbrace{F(\tilde{s},\kappa,0)}_{\neq 0} + D_sF(\tilde{s},\kappa,0)[\,s-\tilde{s}\,] + D_\varepsilon F(\tilde{s},\kappa,0)[\,\varepsilon\,] + R_2(s-\tilde{s},\varepsilon). \tag{42}$$

Even setting $\varepsilon = 0$, the constant term $F(\tilde{s},\kappa,0)$ persists. Any linear approximation based on $\tilde{s}$ must therefore absorb a nonzero residual at order $\mathcal{O}(1)$, so it cannot be first-order correct for the deterministic model. Consequently, its responses to small shocks are biased at first order. Eliminating this bias requires choosing an expansion path $\bar{s}$ satisfying (37).

Remarks. (i) The argument requires only standard smoothness and an invertibility condition on $D_sF(\bar{s},\kappa,0)$. These are the usual conditions underpinning perturbation methods and error bounds (see, e.g., the Implicit Function Theorem and related results as in Stokey–Lucas–Prescott).[16] (ii) Practically, $\bar{s}$ is obtained by solving the deterministic version of the

[16] A short appendix can record these conditions precisely and collect references; here we emphasize the first-order logic fixing the linearization base.

model (with $\varepsilon = 0$) and using that sequence as the expansion path in the piecewise linear algorithm.