

Nonlinear and Stable Perturbation-Based Approximations

Wouter J. Den Haan* and Joris De Wind†

April 26, 2012

Abstract

Users of regular higher-order perturbation approximations can face two problems: policy functions with odd oscillations and simulated data that explode. We propose a perturbation-based approximation that (i) does not have odd shapes, (ii) generates stable time paths, and (iii) avoids the drawbacks that hamper the pruned perturbation approach of Kim, Kim, Schaumburg, and Sims (2008). For models with nontrivial nonlinearities, we find that our alternative and the pruned perturbation approximations give a good qualitative insight in the nonlinear aspects of the true solution, but can differ from the true solution in some quantitative aspects, especially during severe peaks and troughs.

Key Words: accuracy, nonlinear numerical solutions

JEL Classification: C63, E21

*Corresponding author. Department of Economics, London School of Economics and Political Science, Houghton Street, London WC2A 2AE, UK; CEPR, London, UK. e-mail: wjdenhaan@gmail.com.

†CPB Netherlands Bureau for Economic Policy Analysis, Van Stolkweg 14, 2585 JR Den Haag, The Netherlands. e-mail j.de.wind@cpb.nl.

1 Introduction

Perturbation has become a popular choice to solve dynamic stochastic general equilibrium (DSGE) models. Unfortunately, regular higher-order perturbation approximations are not guaranteed to generate non-explosive time paths. Moreover, regular perturbation approximations are polynomials and the unavoidable oscillations of polynomials imply that higher-order approximations do not inherit properties such as monotonicity and convexity from the true underlying policy functions. This is a problem facing all approximation procedures that use polynomials as basis functions. The problem is especially severe for perturbation approximations, because perturbation analysis does not give the user the tools to relocate these problems to those areas of the state space that are of no importance.¹ Consequently, the undesirable oscillations could occur close to the steady state.

To understand the problem, consider the following policy function:

$$x = f(x_{-1}) = -\alpha_0 + x_{-1} + \alpha_1 e^{-\alpha_2 x_{-1}}, \quad (1)$$

with $\alpha_0 = 0.9e^{-1}$, $\alpha_1 = 0.9$, and $\alpha_2 = 1$. The true policy function has a unique fixed point (at $x = 1$) and the dynamics are globally stable. Figure 1 plots this policy function and the second-order perturbation approximation.²

At the fixed point, the second-order perturbation approximation inherits three key properties of the true policy function: (i) increasing in x_{-1} , (ii) strictly convex in x_{-1} , and (iii) the approximation is locally stable, that is, $(\partial f(x_{-1})/\partial x_{-1})|_{x=1} < 1$. For *any* second-order polynomial with these properties, it must be true that the function value goes to $+\infty$ as x_{-1} goes to $+\infty$. This means that the second-order perturbation approximation must have a second intersection with the 45° line, which in turn implies that the dynamics of the approximation are not globally stable.

¹In contrast, the user of projection methods does have this type of control by choosing the appropriate grid. Typically, undesirable oscillations occur outside the grid, which means that one can push these oscillations out of the relevant area by widening the grid. Moreover, uniform convergence is guaranteed if one uses Chebyshev nodes; see Chapter 6.5 in Judd (1998) for a discussion.

²When perturbation analysis is applied to DSGE models, then the derivatives of the unknown policy function are only implicitly defined. In this example, we know the (derivatives of the) policy function and the perturbation approximation is simply the Taylor-series expansion of $f(x_{-1})$.

For the policy function defined in equation 1, the location of the second intersection with the 45° line moves towards the steady state as α_2 increases. If the second intersection is sufficiently far away from the steady state, then the instability will have no practical consequences. In a non-stochastic environment, instability would then only occur when the initial value for x_{-1} is far away from the steady state. In a stochastic environment, the problematic part of the state space would only be reached in the case of extremely unlikely events. In this paper, we will document, however, that these types of problems cannot be ignored in practice.

Kim, Kim, Schaumburg, and Sims (2008) and Lombardo (2010) propose to use pruned perturbation to deal with the problem of exploding simulated data.³ Pruning is already used in several papers.⁴ The pruning procedure does not alleviate the problem that higher-order perturbation approximations can have undesirable odd shapes. Moreover, pruned perturbation approximations have some additional unattractive features. The regular n^{th} -order perturbation approximation uses *one* n^{th} -order policy function to generate *one* time path. In contrast, the pruning procedure generates *multiple* time paths; using the time paths generated by lower-order approximations as the inputs for the higher-order terms in the higher-order approximations. Consequently, the pruning procedure introduces additional state variables. This implies that the pruned perturbation approximation is no longer a function of the original set of state variables. Another striking feature of the pruned perturbation procedure is that the n^{th} -order pruned perturbation approximation does not deliver an exact fit if the truth is an n^{th} -order polynomial even though pruned perturbation approximations are polynomials.⁵ This questions the suitability of pruned perturbation approximations when the underlying function is close to a polynomial.

We propose an alternative, the *perturbation-plus* approximation, which generates stable time paths, does not generate policy functions with odd shapes, and avoids the problems

³In this paper, we propose a slight modification of the procedure of Kim, Kim, Schaumburg, and Sims (2008) to ensure that the way in which higher-order perturbation adjusts the coefficients of the perturbation approximation to the amount of uncertainty is not affected by the pruning procedure.

⁴See, e.g., Fahr and Smets (2010), Doh (2011), Fernández-Villaverde, Guerrón-Quintana, Rubio-Ramírez, and Uribe (2011), and Andreasen (2012).

⁵See online appendix A.2.

of pruning. It starts out with a first-order perturbation approximation. To solve for the period- t model outcomes, we use the *exact* equations of the model for J periods, namely for period t and if $J > 1$, also the $J - 1$ subsequent periods. To obtain a system with as many unknowns as equations, we stipulate that the behavior in period $t + J$ is determined by the first-order perturbation approximation. Although the procedure is easy to program, it is computing intensive unless the chosen value for J is low.

To evaluate whether the pruning and the perturbation-plus approximations are accurate, we consider (i) models where the parameter values are above the critical levels for which regular second-order perturbation approximations generate stable time paths and (ii) models for which regular second-order perturbation generates time paths that reach that part of the state space where the second-order approximation of a monotone increasing function is decreasing. The fact that regular second-order perturbation approximations face difficulties indicates that these are models with nontrivial nonlinearities.⁶ There are many models for which the true solution is close to being linear. If that is the case, then a first-order perturbation approximation is likely to be accurate. If one would use regular higher-order perturbation to solve such models, then the undesirable features highlighted here would most likely occur outside the ergodic set, so one would not have to worry about them.

Neither the pruning nor the perturbation-plus approximations are in general very accurate, although we found some cases where the perturbation-plus approximation is accurate. Although the two modifications to the standard perturbation procedure do not always pass the accuracy tests with flying colors, the news is not all bad. Both procedures provide a good *qualitative* insight in how the true time path differs from the path generated by first-order perturbation. That is, the methods provide a reasonable idea about the nonlinear effects, which are substantial in the models considered. Of course, there is no guarantee that these results carry over to other models and the user should be careful in using these procedures, especially in those cases when the nonlinearities are so important that regular

⁶Although these are nontrivial models, it is straightforward to obtain accurate approximations using projection methods, because the number of state variables is small.

higher-order perturbation approximations generate exploding series.

The organization of this paper is as follows. In section 2, we describe the economic models. In section 3, we explain the problems of higher-order perturbation. In section 4, we discuss the pruning procedure and its drawbacks as well as our alternative, the perturbation-plus approximation. In section 5, we evaluate the accuracy of these two perturbation-based approximations. The last section concludes.

2 Models

In this section, we describe the models. The first model is the neoclassical growth model. The second model is also a very simple model, namely a representative-agent business cycle model in which the labor market is modeled using the Pissarides matching framework. The third model is a simple model in which an agent faces idiosyncratic risk and uses one-period bonds to smooth consumption.

2.1 Neoclassical growth model

The representative agent maximizes

$$\max_{\{c_t, k_t\}_{t=1}^{\infty}} \mathbb{E}_1 \sum_{t=1}^{\infty} \beta^{t-1} \frac{c_t^{1-\gamma} - 1}{1-\gamma}$$

s.t.

$$c_t + k_t = e^{z_t} k_{t-1}^{\alpha} + (1 - \delta)k_{t-1}, \quad (2)$$

$$z_t = \rho_z z_{t-1} + \varepsilon_t, \quad \varepsilon_t \sim N(0, \sigma_z^2), \quad (3)$$

k_0, z_1 given.

The Euler equation is given by

$$1 = \mathbb{E}_t \left[\left(\frac{c_{t+1}}{c_t} \right)^{-\gamma} (\alpha e^{z_{t+1}} k_t^{\alpha-1} + (1 - \delta)) \right]. \quad (4)$$

If $\gamma = \delta = 1$, then the model reduces to the Brock-Mirman model. This model is unusual in the sense that there are analytical solutions for the two policy functions. They are given

by

$$k_t = \alpha\beta e^{z_t} k_{t-1}^\alpha \text{ and} \quad (5)$$

$$c_t = (1 - \alpha\beta) e^{z_t} k_{t-1}^\alpha. \quad (6)$$

2.2 Matching model

There are two types of agents in the model: workers and entrepreneurs. Both types of agents are members of a representative household. At the end of the period, the household receives wages and firm profits from its members. These are distributed among the household's members for consumption.

Firms. In this model, the key decision is made by a representative entrepreneur. The entrepreneur maximizes the discounted value of future firm profits. That is,

$$\max_{\{v_t, n_t\}_{t=1}^{\infty}} \mathbb{E}_1 \sum_{t=1}^{\infty} \beta^{t-1} \left(\frac{c_t}{c_1} \right)^{-\gamma} ((e^{z_t} - w) n_{t-1} - \psi v_t)$$

s.t.

$$n_t = (1 - \rho_n) n_{t-1} + p_{f,t} v_t, \quad (7)$$

$$z_{t+1} = \begin{cases} z_t & \text{with probability } \tilde{\rho}_z \\ -z_t & \text{with probability } (1 - \tilde{\rho}_z) \end{cases} \quad (8)$$

$$n_0, z_1 \text{ given.} \quad (9)$$

Here, v_t is the amount of vacancies posted by the firm, ψ the cost of posting a vacancy, $p_{f,t}$ is the number of matches per vacancy, $p_{f,t} v_t$ is the total number of new hires, ρ_n is the exogenous separation rate, and c_t is the consumption level of the representative household. The wage rate, w , is assumed to be fixed.⁷ The firm takes the value of $p_{f,t}$ as given. Each worker produces e^{z_t} . The value of z_t can take on two values, namely $-\zeta$ and $+\zeta$.

⁷Sticky wages are helpful in generating sufficient employment volatility. Our fixed wage rate is such that the employer never wants to fire the worker and the worker never wants to quit. Our wage rule is, thus, a simple case of Hall (2005).

The first-order conditions are given by

$$\psi = p_{f,t} \lambda_t \text{ and} \tag{10}$$

$$\lambda_t = \beta \mathbb{E}_t \left[\left(\frac{c_{t+1}}{c_t} \right)^{-\gamma} (e^{z_{t+1}} - w + (1 - \rho_n) \lambda_{t+1}) \right], \tag{11}$$

where λ_t is the Lagrange multiplier of the constraint that describes the law of motion for n_t . It represents the value that is generated when an extra worker is added to the firm's workforce.

Consumers. The representative household simply consumes the income earned by its members. Thus,

$$c_t = w n_{t-1} + (e^{z_t} n_{t-1} - w n_{t-1} - \psi v_t) = e^{z_t} n_{t-1} - \psi v_t. \tag{12}$$

Matching market and equilibrium. The number of new hires is determined on a matching market at which firms and the $1 - n_{t-1}$ workers that are not employed search for a match. The total number of matches, m_t , is given by

$$m_t = \phi_0 (1 - n_{t-1})^\phi v_t^{1-\phi}, \tag{13}$$

which means that the number of matches per vacancy is given by

$$p_{f,t} = \phi_0 \left(\frac{1 - n_{t-1}}{v_t} \right)^\phi. \tag{14}$$

Equations (7), (10), (11), (12), and (14) form a system of five equations per period that determine n_t , v_t , λ_t , c_t , and $p_{f,t}$ as a function of n_{t-1} and z_t .

Keeping the problem smooth. The variable $p_{f,t}$ is typically interpreted as a matching probability and is restricted to be less than or equal to 1. This implies that the policy function is no longer smooth. The reason is the following. When z_t takes on very low values, then $\lambda_t < \psi$. If $\lambda_t < \psi$, then the value of an extra employee is less than the posting cost. If $p_{f,t}$ is restricted to be less than 1, then it is impossible to satisfy equation (10). In itself this is not a problem. It simply means that firms post no vacancies, that is, the

firm is at a corner solution. Perturbation analysis can no longer be used, however, if firms occasionally hit corners.

To avoid this dilemma, we do *not* interpret $p_{f,t}$ as a probability and we allow $p_{f,t}$ to exceed 1. That is, if $p_{f,t}$ exceeds 1, then $m_t > v_t$ and firms simply hire more than one worker on each posted vacancy.⁸ If $p_{f,t}$ is allowed to exceed 1, then there always is an internal solution for v_t as long as

$$\lambda_t > 0, \tag{15}$$

a condition that is satisfied in our calibrated model.⁹

2.3 Modified Deaton model

The third model considered is a simple partial equilibrium model in which agents face idiosyncratic income risk. The agent solves the following optimization problem:

$$\max_{\{c_t, a_t\}_{t=1}^{\infty}} \mathbb{E}_1 \sum_{t=1}^{\infty} \beta^{t-1} \left(\frac{c_t^{1-\gamma} - 1}{1-\gamma} - P(a_t) \right)$$

s.t.

$$c_t + \frac{a_t}{1+r} = a_{t-1} + e^{z_t}, \tag{16}$$

$$z_t = \bar{z} + \varepsilon_t \text{ and } \varepsilon_t \sim N(0, \sigma_z^2), \tag{17}$$

$$a_0 \text{ given.} \tag{18}$$

Here, c_t stands for the agent's consumption level, a_t stands for the amount of assets chosen in period t , z_t is an exogenous random income component, and r is the exogenous interest rate. There are two reasons why markets are not complete. First, there is only one financial asset, namely a risk-free bond. Second, there are transactions costs associated with trading in the financial asset, which we model as utility costs. Cash on hand is equal to asset holdings plus income, $x_t = a_{t-1} + e^{z_t}$. Since z_t is assumed to be i.i.d., x_t is the only state variable.

⁸This follows directly from equations (10) and (14).

⁹The restriction that n_t cannot exceed 1 also introduces a nondifferentiability, but this turns out to be irrelevant in our numerical analysis.

The only difference with the model in Deaton (1991) is that we have a penalty function and not a non-negativity constraint. That is, Deaton (1991) assumes that

$$a_t \geq 0. \tag{19}$$

We specify our transactions cost or penalty function such that this inequality constraint is a special case of the model. In particular, the penalty function, $P(a_t)$, is given by

$$P(a_t) = \frac{\eta_1}{\eta_0} \exp(-\eta_0 a_t) + \eta_2 a_t. \tag{20}$$

The value of η_0 controls the curvature of the penalty function and the nonlinearity of the problem. Suppose that $\eta_2 = 0$.¹⁰ Then

$$\lim_{\eta_0 \rightarrow \infty} P(a_t) = \begin{cases} \infty & \text{for } a_t < 0 \\ 0 & \text{for } a_t \geq 0 \end{cases}. \tag{21}$$

That is, as $\eta_0 \rightarrow \infty$ our penalty function implements the standard non-negativity constraint on a_t .

The penalty function not only makes the model more general, it also makes it possible to use perturbation techniques, since perturbation techniques cannot deal with the nondifferentiability introduced by the inequality constraint.¹¹

The Euler equation is given by

$$\frac{c_t^{-\gamma}}{1+r} + \frac{\partial P(a_t)}{\partial a_t} = \beta \mathbb{E}_t [c_{t+1}^{-\gamma}]. \tag{22}$$

¹⁰The term $\eta_2 a_t$ gives additional flexibility, which we exploit in the calibration. The additional term also makes it possible to ensure that the penalty term is equal to zero in the steady state, which may be convenient in some applications.

¹¹De Wind (2008) compares the properties of the model with the inequality constraint given in equation (19) with the corresponding properties of the model with the penalty function for different values of η_0 . Not surprisingly, the value of η_0 matters a lot for the tails of the generated distribution of a_t . In particular, the tail of the distribution generated with the model with a penalty function only matches the tail of the model with the non-negativity constraint for sufficiently high values of η_0 . But model properties such as the volatility of individual consumption depend a lot less on the value of η_0 .

3 Higher-order perturbation in practice

In the introduction, we mentioned two potential problems of higher-order perturbation, namely undesirable shapes and instability. The question arises whether these problems matter, that is, whether they occur in the relevant part of the state space in practice. Of course, the problems do not occur if the shocks hitting the system are *sufficiently* small.¹²

In this section, we shed light on this question by analyzing the perturbation approximations of the solutions to the models of section 2. Before discussing the results, we describe the characteristics of perturbation approximations.

3.1 Characteristics of perturbation approximations

DSGE models can typically be written as follows:

$$0 = E_t [H(x_t, x_{t-1}, y_{t+1}, y_t, z_{t+1}, z_t)], \quad (23)$$

$$z_{t+1} = \Lambda z_t + \varepsilon_{t+1}, \quad (24)$$

$$\varepsilon_{t+1} \sim N(0, \sigma\Omega). \quad (25)$$

Here, x_{t-1} is an $n_x \times 1$ vector containing the state variables of the system, y_t is an $n_y \times 1$ vector containing the endogenous variables that are not state variables, z_t is an $n_z \times 1$ vector with the exogenous random variables, ε_t is the vector with the corresponding innovations, σ is a scalar that controls the overall volatility of the model ($\sigma \geq 0$), and $\sigma\Omega$ is the $n_z \times n_z$ covariance matrix of the innovations. Finally, $H(\cdot)$ is a known vector-valued function with dimension $(n_x + n_y) \times 1$.

We denote the true rational expectations solution by

$$\begin{bmatrix} x_t \\ y_t \end{bmatrix} = \begin{bmatrix} f(x_{t-1}, z_t; \sigma) \\ g(x_{t-1}, z_t; \sigma) \end{bmatrix} \quad (26)$$

and the n^{th} -order perturbation approximation by

$$\begin{bmatrix} x_t \\ y_t \end{bmatrix} = \begin{bmatrix} \tilde{f}_{n^{\text{th}}}(x_{t-1}, z_t; \sigma) \\ \tilde{g}_{n^{\text{th}}}(x_{t-1}, z_t; \sigma) \end{bmatrix}. \quad (27)$$

¹²In fact, linear approximations are accurate if the shocks are "sufficiently" small and the solution is differentiable.

Let s_t denote the arguments of the policy function, that is, $s_t = [x_{t-1}, z_t; \sigma]$, let n_s be the dimension of s , and let \bar{s} denote the corresponding steady state values, that is, $\bar{s} = [\bar{x}, 0; 0]$.¹³ The numerical approximation $\tilde{f}(\cdot)_{n^{\text{th}}}$ is an n^{th} -order perturbation approximation if the following conditions hold:

$$\begin{aligned}
f(s_t)|_{s_t=\bar{s}} &= \tilde{f}_{n^{\text{th}}}(s_t)|_{s_t=\bar{s}} \\
\frac{\partial^i f(s_t)}{\prod_{j=1}^{n_s} \partial s_{j,t}^{i^*(j)}} \Big|_{s_t=\bar{s}} &= \frac{\partial^i \tilde{f}_{n^{\text{th}}}(s_t)}{\prod_{j=1}^{n_s} \partial s_{j,t}^{i^*(j)}} \Big|_{s_t=\bar{s}} \\
\text{and } \forall i &= 1, \dots, n; \quad 0 \leq i^*(j) \leq i \\
&\ni i = \sum_{j=1}^{n_s} i^*(j)
\end{aligned} \tag{28}$$

A similar set of conditions determines whether $\tilde{g}_{n^{\text{th}}}(\cdot)$ is an n^{th} -order approximation.

The true policy functions and the numerical approximations are functions of the endogenous state variables, x_{t-1} , and the exogenous state variables, z_t . This property is an important aspect of recursive models. Our formulations of both the true solution and the regular perturbation approximation allow for the possibility that a particular state variable has no effect on a particular choice, but it does not allow for variables other than the state variables to have an effect. Although this is a standard property of rational expectations solutions and numerical approximations, we highlight this property because, as shown below, the pruning approximation does not satisfy this property.

3.2 Perturbation approximations and the neoclassical growth model

It is well known that both log-linear and linear approximations are accurate for the neoclassical growth model as long as σ_z takes on plausible values. In fact, the solution of the Brock-Mirman model is a log-linear function of capital and productivity. Consequently, *any* perturbation approximation will recover the true rational expectations solution if the model is written in the logarithms of the variables. To ensure that the problem remains

¹³The perturbation approximation procedure is a Taylor series expansion around the steady *and* around the case with $\sigma = 0$. This is the reason why we explicitly include σ as an argument of the policy function.

nontrivial, even when we consider the Brock-Mirman version, we calculate approximations in the levels of the state variables, not in the logarithms.

Consistent with the facts stated in the last paragraph, we find that the higher-order terms of higher-order perturbation approximations are quantitatively small for parameter values commonly used in the literature. The perturbation approximations are then not affected by the problems discussed in the introduction. Therefore, we also consider values for σ_z that are higher than those normally used. Although the objective of this section is to analyze the problems of higher-order perturbation approximations researchers encounter in practice, it is still useful to consider these not so typical parameter values. There are two reasons. First, the simplicity of the model makes it easy to understand why higher-order perturbation approximations run into problems; a better understanding of the problems in such a simple case is helpful in understanding the problems in more complex cases. Second, the analysis points out that the highlighted problems eventually show up for some parameter values, *even* in models with very simple nonlinearities.

Non-monotonicity of second-order perturbation approximations. Panel A of Figure 2 plots the perturbation approximation of the capital policy function when σ_z is equal to 0.007, a very standard value. The policy functions are plotted as a function of k_{t-1} for three different values of z_t .¹⁴ The results are shown for the Brock-Mirman model, that is, $\gamma = \delta = 1$. The other parameter values take on standard values.¹⁵

It is impossible for standard second-order perturbation approximations to be monotonically increasing (or decreasing). As documented by the figure, the problematic decreasing part occurs, however, when the capital stock is very high, namely when it is more than 2.5 times its steady state value. Simulated values for the capital stock would not reach such high values when $\sigma_z = 0.007$.¹⁶ If the approximation remains unchanged, then an increase in σ_z would simply increase the volatility of capital, making it more likely that

¹⁴The values are 0 and plus and minus two times the standard deviation of z .

¹⁵In particular, $\alpha = 0.36$, $\beta = 0.99$, and $\rho_z = 0.95$.

¹⁶In a simulation of 10,000 observations the largest (smallest) value for capital is equal to 0.2258 (0.1996) when $\sigma_z = 0.007$.

capital reaches the undesirable part of the approximation. But the second-order perturbation approximation changes when σ_z changes. In particular, it could be the case that the undesirable decreasing part of the policy functions is pushed to the right as σ_z increases. This would make it more difficult to reach this problematic part of the state space. This turns out *not* to be the case, as is documented in Panel B of Figure 2. This panel plots the policy functions when σ_z is equal to 0.2. For the high value of z_t , the turning point after which the function is decreasing in k_{t-1} is indeed pushed to the right. For the low value of z_t , however, the turning point is pushed to the left. In fact, the policy function at the low productivity value is downward sloping for a large *relevant* range of values of k_{t-1} .¹⁷

Instability of second-order perturbation approximations. According to the second-order perturbation approximation, the choice for capital and consumption are strictly positive at $k_{t-1} = 0$. In fact, the chosen levels are quite high. For plausible values of σ_z one would not get close to such low values of capital. But as σ_z increases, the volatility increases and the ergodic set expands. Moreover, as σ_z increases, the policy function for capital shifts down (keeping the value of z fixed) making low values of capital only *more* likely.

Of course, it is not accurate at all to have positive capital and consumption choices when $k_{t-1} = 0$.¹⁸ Although inaccurate, this property does make the perturbation approximation more stable than the true solution. When using the second-order perturbation approximation, one could even start at somewhat *negative* capital levels and the economy will still revert back to the area around the steady state.¹⁹ We find this to be true for

¹⁷In a simulation of 10,000 observations the largest (smallest) value for capital is equal to 6.8675 (0.0069) when $\sigma_z = 0.2$. The standard deviation is equal to 0.44.

¹⁸According to the true rational expectations solution, both capital and consumption should be set equal to 0 when the beginning-of-period capital stock equals 0; this is the only feasible choice.

¹⁹The second-order perturbation solution for k_t does have a second fixed point, but the value of k at this second fixed point is negative. If the capital stock would ever get below this negative value, then the second-order perturbation solution would be explosive. Note that $\lim_{k_{t-1} \rightarrow \infty} \tilde{f}_{2nd}(k_{t-1}, z_t) = -\infty$. That is, if the economy would start out at (or reach) a sufficiently high capital stock, then the capital choice

a wide range of parameter values, also when we move beyond the Brock-Mirman model and consider other values for γ and allow for partial depreciation.²⁰ Nevertheless, the instability problem of the second-order perturbation approach is just around the corner. This is discussed next.

The second-order perturbation approximation grossly violates the budget constraint for low values of the beginning-of-period capital stock. But it is easy to obtain a second-order perturbation approximation that exactly satisfies the budget constraint. In particular, one could use the perturbation approximation for either capital or consumption and use the budget constraint to solve for the other variable. The numerical approximation satisfies the conditions in equation (28) and is, thus, a perturbation approximation.

Figure 3 plots the perturbation approximation for capital, $\tilde{f}_{2\text{nd}}(k_{t-1}, 0)$, when consumption is determined by the standard second-order perturbation approximation and capital is solved from the budget constraint. The value of z is set equal to 0. The policy function then gives the dynamics of the system if there are no shocks.²¹ The solid line corresponds to the case when the parameters are the same as those used to create Panel A of Figure 2.

Interestingly, the perturbation approximation is now monotonically increasing for all positive values of k_{t-1} . Thus, in terms of avoiding odd shapes, this alternative is an improvement. In terms of stability it is not. The figure documents that the policy function for capital has an additional *positive*-valued fixed point to the left of the steady state.^{22,23} could be so negative that the solution gets into the unstable region.

²⁰We could find a case with a second *positive*-valued fixed point, but the value of k at this second fixed point is very small and we had to raise the value of γ to 35. In this case the economy would diverge for positive initial values when these initial values are below this second fixed point.

²¹This policy function is also indicative of the *expected* dynamics if there are shocks.

²²If γ is equal to 1, then this second fixed point of the perturbation approximation occurs at a value of the capital stock that is only 10% of the steady state value. The figure also plots the second-order perturbation approximation for the case when γ is equal to 10. In this case, the value of capital at the second fixed point is substantially higher, namely around 30% of the steady state.

²³In fact, there is a third positive fixed point. But this fixed point is quite far away from the steady state and is ignored in the text. But if the economy would start out at capital levels above this third fixed point, then capital is expected to grow without bound.

The time paths that start out or reach such low capital levels are diverging.

Summary for the neoclassical growth model. In addition to the cases discussed here, we have considered the properties of second-order perturbation approximations for several parameter values including those that allow for partial depreciation. We found that odd shapes and instability only happen in unlikely cases. That is, in practice second-order perturbation approximations of the neoclassical growth model do not exhibit the problems we highlighted in the introduction. Nevertheless, we do not consider the analysis here very comforting. It is true that one has to go outside the usual range of parameter values to encounter problems. But we have shown that higher-order perturbation approximations do have undesirable shapes and instability problems, even though the model is very simple and has a log-linear solution. This brings up the question whether the problems will occur for more standard parameter values when more interesting models are considered. We document in the next two subsections that we have to answer this question in the affirmative.

3.3 Perturbation approximations and the matching model

In the previous section, we documented problematic features of perturbation approximations to the neoclassical growth model. But these problematic features are only relevant when parameter values are such that the generated volatility is much larger than what is observed for aggregate data. In this subsection, we consider the matching model. In contrast to the results for the neoclassical growth model, we will show that the standard second-order perturbation approximation generates explosive time paths when the model is calibrated to generate a *realistic* amount of volatility. As discussed in section 2.2, the productivity level is assumed to be a discrete-valued random variable that can take on only two values. The main advantage of this assumption is that it allows us to portray the reasons behind the problems of higher-order perturbation approximations with a simple graphical analysis. The other advantage of this assumption is that it is easy to ensure that the true model solution is always well-defined. In particular, we assume that the low

value of z_t is such that profits are always positive.²⁴

Parameter values for the matching model. We choose ζ and $\tilde{\rho}_z$ such that the standard deviation of z_t is equal to 0.007 and the autocorrelation of z_t is equal to 0.95.²⁵ These are typical values in the business cycle literature. We set the discount factor, β , equal to 0.99 and the curvature parameter of the matching function, ϕ , equal to 0.5.²⁶ The values of the posting cost, ψ , the separation rate, ρ_n , and the scaling factor in the matching function, ϕ_0 , are chosen such that the steady state values of the unemployment rate, the number of matches per unemployed worker, and the number of matches per vacancy, are equal to 5%, 0.7, and 0.7, respectively.²⁷ The value of γ is set equal to 4.5. This value for the coefficient of relative risk aversion is perhaps a bit higher than the most commonly used values, but still a plausible value.²⁸ We consider two values for the wage rate, namely $w = 0.96$ and $w = 0.973$. When w is equal to 0.96, then the volatility of the employment relative to the volatility of productivity is equal to 0.25, whereas the observed ratio in the data is equal to 0.437.²⁹ Thus, to match the observed relative volatility, the value of w has to be increased. When $w = 0.973$, then the volatility of employment relative to labor productivity (calculated using an accurate projection method) matches its observed counterpart.³⁰

²⁴This condition is sufficient to keep the problem well-defined, but it is not necessary. That is, we could allow profits to be somewhat negative. To keep the problem well-defined it is important that $\lambda_t > 0$. As long as expected future profits offset current losses, then λ_t would be positive.

²⁵This means that $\zeta = 0.0224$ and $\tilde{\rho}_z = 0.975$.

²⁶See Petrongolo and Pissarides (2001) for a motivation for the chosen value of ϕ .

²⁷This implies that $\psi = 0.5965$, $\rho_n = 0.368$, and $\phi_0 = 0.7$.

²⁸For lower values of γ , the second-order perturbation solution is still well-behaved when the volatility of employment relative to the volatility of productivity is equal 0.437, i.e., the observed value. For somewhat higher values of the target, the second-order perturbation approximation runs into the same problems as those discussed here. For example, when $\gamma = 3.5$ and $w = 0.976$, then the relative volatility of employment would be equal to 0.50 according to the accurate projection solution and the 2nd-order perturbation approximation is no longer stable.

²⁹The series are filtered using the HP filter. See Den Haan and Kaltenbrunner (2009) for further details on the data used.

³⁰See online appendix C.1 for a description of the projection method used.

Second-order perturbation when $w = 0.96$. Panel A of Figure 4 plots the second-order perturbation approximation for employment, n_t , when the wage rate is equal to 0.96. It also plots a very accurate solution obtained with a projection method, which we use as a stand-in for the truth.³¹ When z takes on its low value ($-\zeta$), then the perturbation approximation has a second fixed point when n is (roughly) 0.83. With discrete support this second fixed point is irrelevant in the sense that as long as the economy starts in the ergodic region (also indicated in the figure), then the generated time path is always stable. In particular, even if the economy is extremely unlucky and z_t is always equal to $-\zeta$, then the economy will still converge towards a positive value for n_t . Similarly, if z_t is always equal to $+\zeta$, then the generated time path will remain well-behaved.

Moreover, the second-order perturbation approximation is a monotone increasing function in the ergodic set. In fact, the second-order perturbation approximation is very close to the very accurate projection approximation in the ergodic set.

Even though the second-order perturbation approximation is doing very well, a small change in the parameter values changes the picture completely. And the properties of the model ask for such a change in the parameter values. As mentioned above, the standard deviation of HP-filtered employment relative to the standard deviation of HP-filtered productivity is still below its empirical counterpart. In Hagedorn and Manovskii (2008), it is shown that this ratio increases when the average surplus decreases. To match observed employment volatility we increase the value of the wage rate, w , from 0.96 to 0.973.

Second-order perturbation when $w = 0.973$. Panel B of Figure 4 plots the policy function when $w = 0.973$. The small increase in the wage rate leads to a minor shift in the policy function. This is true for the "true" policy function and the second-order perturbation approximation. Nevertheless this minor shift has enormous consequences for time paths simulated with the second-order perturbation approximation because the time path for employment, n_t , now explodes.³² The reason is that the second-order perturbation approximation is now *always* below the 45° line when z takes on its low value. That is,

³¹See online appendix C.1 for details.

³²Instability already occurs for lower values of the wage rate, namely when it is close to 0.965.

according to the perturbation approximation there is no longer a bounded ergodic set.

3.4 Perturbation approximations and the modified Deaton model

The modified Deaton model resembles the first model in two aspects. First, it is also a relatively simple model. Second, the agent in the modified Deaton model also faces a trade-off between the return on savings (or cost of borrowing) and consumption smoothing. The model also differs in an important aspect from the first model. The volatility is much higher at *standard* parameter values, since this is a model that describes individual behavior and incorporates *idiosyncratic* uncertainty.

Parameter values for the modified Deaton model. We set $r = 0.03$, $\gamma = 3$, $\bar{z} = 0.4$, $\sigma_z = 0.1$, and $\beta = 0.9$. We choose a low value for β to ensure that agents are sufficiently impatient. If agents are not impatient, then penalty functions or borrowing constraints would not matter. The standard deviation of idiosyncratic income, σ_z , is such that a negative two-standard-deviations shock implies a level of income that is 20% below its mean value. This is clearly not an excessively volatile process for idiosyncratic risk; some papers in the literature even consider specifications for idiosyncratic risk that allow individual income levels to be equal zero.³³ Furthermore, we set $\eta_0 = 20$, $\eta_1 = 0.04464$, and $\eta_2 = 0.00352$. The values of η_1 and η_2 are such that the mean and standard deviation of a_t are equal to the corresponding values in the model with the non-negativity constraint. The value of η_0 is modest in the sense that 19% of the times the agent chooses a negative value for a_t , whereas this cannot happen in the original Deaton model. To get closer to the Deaton model one would need a higher value for η_0 , that is, more curvature in the penalty function.

Instability of second-order perturbation approximation. Figure 5 plots the "true" policy function and the second-order perturbation approximation.³⁴ Instead of plotting

³³See, for example Krusell and Smith (1998).

³⁴The "true" policy function is represented by a very accurate solution obtained with a projection procedure. See online appendix C.2 for details.

the savings choice, a_t , as a function of the state variable, cash on hand, x_t , we plot the expected value of next-period's cash on hand, $E[x_{t+1}|x_t]$, which is equal to $a_t + E_t[e^{z_{t+1}}]$. The reason is that the graph of the relationship between (the expected value of) x_{t+1} as a function of x_t directly reveals whether the dynamics are stable or not.

As documented in the figure, the true policy function is convex and monotonically increasing for the values of x_t observed in the cross-section. Interestingly, the same is true for the second-order perturbation approximation in the relevant part of the state space. The second-order perturbation approximation is unstable, however, because there is a second intersection with the 45°-line above the true steady state. If this intersection is too close to the true steady state, then simulated time paths will eventually take on values to the right of the second intersection at which point the economy is expected to diverge.³⁵ This turns out to be the case for the parameter values used here.

Instability of higher-order perturbation approximations. The question arises whether one can avoid or at least reduce the severity of the instability of the second-order perturbation approximation by going to higher-order approximations. To shed light on this question, we plot the third, fourth, and fifth-order perturbation approximations for the modified Deaton model in Figure 6.

As predicted by theory, the approximation around the steady state improves if one increases the order of the approximation, although the scale of the figure is such that this is not very clear. For our purpose, the global fit is more interesting. The figure documents that the third-order perturbation approximation, like the second-order approximation, has a second intersection with the 45°-line and that this undesirable intersection is *closer* to the steady state than the undesirable additional intersection for the second-order perturbation approximation. Consistent with this fact, we find that series generated with the third-order perturbation approximation start an explosive trajectory faster than series generated with the second-order perturbation approximation.

³⁵Being to the right of the intersection is not sufficient for the asset holdings of the agent to diverge. If the agent receives sufficiently soon sufficiently low values for z_t , then the asset holdings would drop back into the stable region.

The fourth and fifth-order perturbation approximations are stable, but the stability comes at the cost of having a sharply decreasing policy function in some part of the state space. All polynomial approximations display oscillations and the higher the order of the polynomial the more oscillations. The problem with perturbation is that one cannot control where such undesirable oscillations occur. The histogram below the graph displays the observed distribution of the state variable according to the accurate projection approximation. It documents that the decreasing part occurs in a very relevant part of the state space for the fifth-order approximation.

4 Stable nonlinear perturbation-based approximations

In this section, we describe two procedures—both based on perturbation analysis—that generate stable time paths. The first procedure solves for the period- t decisions using (i) the *exact* nonlinear equations of the model for periods t through $t + J$ and (ii) the first-order perturbation approximation to describe the behavior in period $t + J$.³⁶ The second procedure is perturbation with pruning. The pruning procedure resembles the standard perturbation procedure in that it starts out with a first-order approximation and sequentially adds higher-order monomial terms. In contrast to the standard perturbation approach, the pruning procedure also adds state variables at each step.

4.1 Perturbation-plus procedure

To explain the procedure, it is more convenient to write the generic description of a DSGE model using only the state variables. That is, we replace equation (23) with

$$0 = \mathbf{E} [H(x_{+1}, x, x_{-1}, z_{+1}, z)]. \quad (29)$$

The objective is to determine x given values for x_{-1} and z . There are two reasons why equation (29) is not a standard nonlinear equation in x given x_{-1} and z . First, the equation contains an integral over the realizations of z_{+1} . Second, x_{+1} is also unknown.

³⁶Any numerical approximation could be used. Thus, this procedure can be used to improve any numerical approximation.

Leading equation (29) with one period gives an extra equation, but also introduces an extra variable, namely x_{+2} . The first issue can be dealt with using numerical integration procedures. To deal with the second issue we replace x_{+j} by the first-order perturbation approximation, $\tilde{f}_{1\text{st}}(x_{+j-1}, z_{+j})$. If we use the *one-step* ahead perturbation-plus procedure, then $j = 1$ and x_{+1} is replaced by $\tilde{f}_{1\text{st}}(x, z_{+1})$. Thus, we solve x from

$$0 = \tilde{\mathbb{E}} \left[H(\tilde{f}_{1\text{st}}(x, z_{+1}), x, x_{-1}, z_{+1}, z) \right], \quad (30)$$

where $\tilde{\mathbb{E}}$ denotes that the integral is calculated using a numerical integration procedure. We denote the value of x that solves equation (30) by $\hat{f}_{+1}(x_{-1}, z)$, where the subscript indicates that x is solved using the true model equations for 1 period. There may be no analytical solution for $\hat{f}_{+1}(x_{-1}, z)$ in which case $\hat{f}_{+1}(x_{-1}, z)$ is only implicitly defined by equation (30). In this case, one has to use a nonlinear equation solver to solve for the value of $\hat{f}_{+1}(x_{-1}, z)$.

To generate a time path for x_t one would simply iterate on $\hat{f}_{+1}(x_{t-1}, z_t)$. Note that $\hat{f}_{+1}(x_{t-1}, z_t)$ is constructed under the assumption that next period's value for x is calculated using $\tilde{f}_{1\text{st}}(x_t, z_{t+1})$. As next period comes along, however, x_{t+1} is not calculated using the first-order perturbation approximation, $\tilde{f}_{1\text{st}}(x_t, z_{t+1})$, but is calculated using $\hat{f}_{+1}(x_t, z_{t+1})$.

The two-step ahead modification, $x = \hat{f}_{+2}(x_{-1}, z)$, is the value of x that is the solution to

$$0 = \tilde{\mathbb{E}} \left[H(\hat{f}_{+1}(x, z_{+1}), x, x_{-1}, z_{+1}, z) \right]. \quad (31)$$

That is, the value of x is based on (i) the exact equations of the model for this period (ii) the exact equations of the model for the next period, and (iii) the assumption that the behavior in the period after the next period is based on the first-order perturbation approximation.³⁷

³⁷The perturbation-plus procedure is related to the extended-path method of Fair and Taylor (1983) and Gagnon (1990). This procedure also solves for period- t variables by looking a number of periods into the future. The difference is that the perturbation-plus procedure calculates conditional expectations explicitly, namely by using numerical integration procedures, and it closes the system by using the first-order perturbation solution to describe the *behavior* in the last period, whereas the extended-path method

In theory, one could iterate on this process and construct the J -step ahead modification, $x = \widehat{f}_{+J}(x_{-1}, z)$. As J approaches infinity, then x and x_{+1} are based on the same policy function, that is, one has a rational expectations equilibrium.

The procedure quickly becomes expensive as J increases, especially when there is no analytical expression for $\widehat{f}_{+J}(x_{-1}, z)$. Even in a model as simple as the neoclassical growth model, there is no analytical expression for the value of x that solves equation (30). As discussed in online appendix B.1, there is a slight modification of the algorithm that is much faster and we found the numerical results to be very similar.³⁸

Discussion. The perturbation-plus procedure is quite easy to program. Relative to implementing a perturbation procedure, no new tools are needed except a numerical integration (quadrature) procedure. But this is not difficult; using a quadrature technique is like programming the expectation of a random variable with discrete support. The perturbation-plus procedure is easier to implement than projection methods because it does not require the user to construct a grid. The user also does not have to worry about what class of approximating functions to use (Chebyshev polynomials, splines, etc.). Nevertheless, the procedure has several of the benefits of a projection procedure, at least for sufficiently high J , because it uses the exact equations of the model and it explicitly approximates the conditional expectation with an accurate numerical integration procedure.

The disadvantage of this procedure is that the problem quickly gets very expensive as J increases. That is, in practice one can only use this procedure if the appropriate value for J is low. Whether an accurate solution can be obtained with a low value for J depends on the model at hand and in particular on the value of β . In section 5, we discuss an example in which the first-order as well as the second-order perturbation approximations are not accurate, but the one-step ahead perturbation-plus approximation is accurate.

closes the system using terminal conditions for expectations or variables.

³⁸Our finding that the results are similar could very well depend on the problem at hand and not carry over to other models.

4.2 Pruning

In this section, we describe the pruning procedure. Our procedure is based on Kim, Kim, Schaumburg, and Sims (2008), but—as discussed below—we deviate from their procedure in one small aspect.³⁹ We follow Collard and Juillard (2001) and define the stochastic steady state as the fixed point of the regular perturbation solution—*with* a possible correction for uncertainty—when z_t is equal to its steady state value. That is, $\bar{x}_{2\text{nd}}(\sigma)$ is the solution to^{40,41}

$$\bar{x}_{2\text{nd}} = \tilde{f}_{2\text{nd}}(\bar{x}_{2\text{nd}}, \bar{z}; \sigma). \quad (32)$$

From now on we subtract the non-stochastic steady state from the state variables in the formulation of the perturbation approximation. To avoid introducing new notation, we will denote this approximation also with $\tilde{f}_{2\text{nd}}(\cdot)$.

This second-order approximation can always be written as follows:⁴²

$$\begin{aligned} \tilde{f}_{2\text{nd}}(x_{t-1} - \bar{x}_{2\text{nd}}, z_t - \bar{z}; \sigma) \\ = \\ \tilde{f}_{2\text{nd}}^{(1)}(x_{t-1} - \bar{x}_{2\text{nd}}, z_t - \bar{z}; \sigma) + \tilde{f}_{2\text{nd}}^{(2)}(x_{t-1} - \bar{x}_{2\text{nd}}, z_t - \bar{z}; \sigma), \end{aligned} \quad (33)$$

where $\tilde{f}_{2\text{nd}}^{(1)}(\cdot)$ is the part of $\tilde{f}_{2\text{nd}}(\cdot)$ with the linear terms, and $\tilde{f}_{2\text{nd}}^{(2)}(\cdot)$ is the part of $\tilde{f}_{2\text{nd}}(\cdot)$ with the second-order terms.

The pruning procedure consists of the following steps.

1. Simulate $x_t^{(1)}$ using

$$x_t^{(1)} - \bar{x}_{2\text{nd}} = \tilde{f}_{2\text{nd}}^{(1)}(x_{t-1}^{(1)} - \bar{x}_{2\text{nd}}, z_t - \bar{z}; \sigma). \quad (34)$$

³⁹Online appendix A gives a more detailed discussion and also discusses the differences with the procedure of Lombardo (2010).

⁴⁰In the remainder of the paper we will write $\bar{x}_{2\text{nd}}$ instead of $\bar{x}_{2\text{nd}}(\sigma)$, but the reader should remember that a steady state indexed by an i^{th} subscript indicates a stochastic steady state, which depends on σ .

⁴¹The non-stochastic steady state is the fixed point when $\sigma = 0$ using the true rational expectations solution, i.e., the value of x that solves $x = f(x, \bar{z}; 0)$, which is the same as the value for x that solves $H(x, x, x, \bar{z}, \bar{z}) = 0$.

⁴²For second-order perturbation, the value of σ at most affects the stochastic steady state. So strictly speaking we would not need σ as a separate argument for any of the functions in equation (33). For higher-order approximations, σ is needed as a separate argument.

2. Simulate $x_t = x_t^{(2)}$ using

$$x_t^{(2)} - \bar{x}_{2^{\text{nd}}} = \tilde{f}_{2^{\text{nd}}}^{(1)}(x_{t-1}^{(2)} - \bar{x}_{2^{\text{nd}}}, z_t - \bar{z}; \sigma) + \tilde{f}_{2^{\text{nd}}}^{(2)}(x_{t-1}^{(1)} - \bar{x}_{2^{\text{nd}}}, z_t - \bar{z}; \sigma), \quad (35)$$

where the values of z_t used are identical to those used in step 1. The process $\tilde{f}_{2^{\text{nd}}}^{(2)}(x_{t-1}^{(1)} - \bar{x}_{2^{\text{nd}}}, z_t - \bar{z}; \sigma)$ is stationary because both $x_t^{(1)}$ and z_t are stationary. The process $\tilde{f}_{2^{\text{nd}}}^{(1)}(x_{t-1}^{(2)} - \bar{x}_{2^{\text{nd}}}, z_t - \bar{z}; \sigma)$ is stationary unless the Blanchard-Kahn conditions are not satisfied. Consequently, the simulated values of x_t are stationary as well.

The n^{th} -order *pruned* perturbation approximation, $x_t = x_t^{(n)}$, is generated using the following iterative scheme:

$$x_t^{(j)} - \bar{x}_{n^{\text{th}}} = \sum_{i=1}^j \tilde{f}_{n^{\text{th}}}^{(i)}(x_{t-1}^{(j-i+1)} - \bar{x}_{n^{\text{th}}}, z_t - \bar{z}; \sigma) \text{ for } j = 1, \dots, n. \quad (36)$$

For $n \geq 2$, there is an alternative. First calculate $x_t^{(1)}$ and $x_t^{(2)}$ using equation (36). Next calculate the n^{th} -order approximation using

$$x_t^{(j)} - \bar{x}_{n^{\text{th}}} = \tilde{f}_{n^{\text{th}}}^{(1)}(x_{t-1}^{(j)} - \bar{x}_{n^{\text{th}}}, z_t - \bar{z}; \sigma) + \sum_{i=2}^j \tilde{f}_{n^{\text{th}}}^{(i)}(x_{t-1}^{(j-i)} - \bar{x}_{n^{\text{th}}}, z_t - \bar{z}; \sigma) \quad \text{for } j = 3, \dots, n. \quad (37)$$

The formulation in equation (37) uses the highest available (stationary) variable. Even though the specification in equation (36) does not use this possible update, it is also a n^{th} -order approximation since $x_t^{(j)}$ is used in the $(n - j + 1)^{\text{th}}$ -order monomials.

There is a small difference between the procedure described here and the way pruning is implemented in Kim, Kim, Schaumburg, and Sims (2008). In each step of the n^{th} -order pruned perturbation procedure, we use the coefficients of the n^{th} -order perturbation approximation. For example, to generate $x_t^{(1)}$ we use the linear part of the n^{th} -order perturbation approximation whereas Kim, Kim, Schaumburg, and Sims (2008) use the first-order perturbation approximation. As is documented in online appendix A, both approximations are n^{th} -order approximations.

The motivation for this modification is straightforward. One reason to use higher-order perturbation approximations is that they allow uncertainty to affect the coefficients

of the approximation. For example, the constant term in the second-order perturbation approximation depends in general on the amount of uncertainty.⁴³ It is this dependence of the constant term on uncertainty that allows the second-order perturbation approximation to capture the impact of uncertainty on, for example, average savings behavior. Our version of the pruning procedure ensures that this dependence is not lost in the pruning procedure. For example, the stochastic steady state of our pruned perturbation solution is identical to the stochastic steady state of the underlying regular perturbation solution. This is not the case for the pruning procedure of Kim, Kim, Schaumburg, and Sims (2008).^{44,45}

The pruned perturbation approximation is a recursive function just like the regular perturbation approximation, but of a larger (and potentially much larger) set of variables. It is a *correspondence* in the original state variables, x_{-1} and z . To illustrate this aspect of the pruned perturbation approximation, we use the neoclassical growth model. We assume that the law of motion for productivity, z_t , is equal to the first-order Markov process described in section 2.2. The advantage of this process is that z_t takes on only two values, which makes it possible to graphically document whether the policy function is close to a function of the original state variables or not.

To see whether the pruned perturbation approximation is close to a function of the original state variables we do the following. Using the pruned second-order perturbation approximation, we generate a long time series for capital, $k_t^{(2)}$.⁴⁶ Next, we plot the change

⁴³The stochastic steady state of the perturbation approximation is not equal to the non-stochastic steady state exactly because of this dependence.

⁴⁴In generating $x_t^{(1)}$, Kim, Kim, Schaumburg, and Sims (2008) use the first-order perturbation solution, not the linear part of the higher-order perturbation solution. Since $x_t^{(1)}$ is simply the first-order perturbation solution, its stochastic steady state is the non-stochastic steady state, not the stochastic steady state of the underlying perturbation solution. Since $x_t^{(1)}$ is used as input in subsequent steps, the stochastic steady state of series generated in subsequent steps is also affected.

⁴⁵As discussed in online appendix A, both procedures generate valid n^{th} -order perturbation approximations. The reason is these are *local* approximations, i.e., when $\sigma \rightarrow 0$, and the stochastic steady state and the non-stochastic steady state converge towards each other as $\sigma \rightarrow 0$.

⁴⁶The simulation is based on the following parameter values: $\alpha = 0.36$, $\beta = 0.99$, $\gamma = 3$, $\delta = 0.025$, $\tilde{\rho}_z = 0.975$, and $\zeta = 0.6405$.

in capital, $k_t^{(2)} - k_{t-1}^{(2)}$, as a function of the state variable, $k_{t-1}^{(2)}$. One set of numbers will be for the high value of z_t and one set for the low value of z_t . The results are reported in Figure 7. The figure clearly illustrates that the pruned perturbation approximation does not come close to being a function of the original state variables.⁴⁷

5 Accuracy of the stable perturbation-based procedures

In this section, we evaluate the accuracy of the pruned perturbation and the perturbation-plus approximations. For the Brock-Mirman model, we compare a time path of the approximation with the corresponding time path of the true solution. For the other models, we use a time path generated by a very accurate projection method instead of the (unknown) true solution.⁴⁸ The length of the time path, T , is set equal to 10,000.

The distance between a period- t variable generated with an approximation and the corresponding "true" value could be measured as the absolute percentage error. This will not make sense if—as is the case in several of our models—variables take on values that are close to zero. An alternative would be to scale the (absolute) difference by the time series average of the true series. But this measure would overestimate the seriousness of deviations if variables take on values that are bigger than multiple times this mean and we encounter observations where the true value is more than 20 times the mean. For those type of observations, the regular percentage error would be more appropriate.

Therefore, we define the period- t error as

$$e_t = \min \left\{ \left| \frac{\tilde{x}_t - x_t}{x_t} \right|, \left| \frac{\tilde{x}_t - x_t}{\bar{x}_T} \right| \right\}, \quad (38)$$

where \tilde{x}_t is the realization according to the approximation, x_t the realization according to

⁴⁷The value of ζ used to generate the graph is equal to 0.6405, which means that the amount of uncertainty is high for a macroeconomic model. But it does not make sense to analyze this issue for standard parameter values, since for standard parameter values the solution to this very simple model is approximated well with first-order perturbation and there would be no reason to consider higher-order pruned perturbation. Den Haan and De Wind (2009) illustrate this aspect of pruned perturbation approximations using the modified Deaton model.

⁴⁸That the projection method generates very accurate solutions is established in online appendix C.

the true solution (or very accurate projection method solution), and \bar{x}_T is the mean value of x_t . We report both the maximum and the mean of e_t . For all three models, we report only the results for the state variable. The errors for consumption are somewhat smaller, but the conclusions drawn here do not depend on the variable considered.

5.1 Accuracy of approximations to the Brock-Mirman model

We consider two values for σ_z , namely $\sigma_z = 0.1$ and $\sigma_z = 0.2$. Table 1 reports information about the approximation errors and table 2 reports summary statistics for the behavior of capital according to the different policy rules.⁴⁹ The following observations can be made.

First, the first-order perturbation approximation performs very poorly. This approximation generates large *negative* values for the capital stock. Moreover, the highest observation for capital that is generated by the first-order perturbation approximations is more than 50% below the correct highest observation value when $\sigma_z = 0.1$ and more than 87% below the correct peak when $\sigma_z = 0.2$. Not surprisingly, the accuracy measures are very poor at these high values for σ_z . The first-order approximation would do fine for values of σ_z that are typical for representative-agent models. The point here is to show that even models with simple nonlinearities eventually get into serious problems if the amount of uncertainty is increased.

Second, the pruned second-order approximation performs worse than the regular second-order perturbation approximation. Although, the pruned second-order approximation typically does substantially better than first-order perturbation, its maximum error is higher when $\sigma_z = 0.2$. Whereas the problem of first-order perturbation is that it generates values for capital that are way too low, the problem for both the regular and the pruned second-order perturbation approximation is that the lowest values are substantially above the true minimum value. The gap is roughly equal to one quarter of the mean capital

⁴⁹The minimum values obtained by the second-order perturbation approximations are the same when $\sigma = 0.1$ and when $\sigma = 0.2$. The reason is the following. First, the correction term for uncertainty is almost zero. Second, when z is low then the (quadratic) policy function is (i) quite flat and (ii) close to the minimum. This means that the chosen values of k hover around the same value until a sufficiently high value for z pushes the economy out of this area.

stock.

Third, our proposed alternative does substantially better than the second-order perturbation approximations, both in terms of having better accuracy measures and in terms of generating summary statistics that are closer to the truth. For the two-step ahead perturbation-plus approximation, the maximum errors are equal to 5.2% and 8.1% when σ_z is equal to 0.1 and 0.2, respectively. Although nontrivial numbers, they are substantially lower than for the second-order pruned approximation for which the corresponding errors are 47.9% and 193.8%.

Figure 8 plots that part of the sample where the largest errors are obtained by the perturbation-plus approximations when $\sigma_z = 0.2$. This truly is an usual period: the true value of k_t takes on a value that is more than twenty times the sample average. The two-step ahead perturbation-plus approximation does a good job following the true time path, but does not reach the same peak. It reaches a maximum value of 6.3 whereas the true maximum is equal to 6.9. The maximum reached by second-order pruned perturbation is only 2.2. The three-step ahead perturbation-plus approximation does even better than the two-step ahead approximation. For example, the maximum reached during this enormous upswing is equal to 6.68, only 2.7% below the true maximum (which is 20 times above the average value). The problem is that the three-step ahead perturbation-plus approximation is quite expensive to run.⁵⁰

The figure documents that the first-order perturbation solution grossly and persistently underestimates the upswing of capital during this period. Since this approximation is used to describe future behavior in the perturbation-plus procedure it is not surprising that the perturbation plus procedure also underpredicts capital. In fact, it is quite amazing that most of the gap between the true solution and the first-order perturbation solution can be corrected by using the actual equations of the model for only two periods and using this

⁵⁰With 5 quadrature nodes it takes roughly 0.021 seconds per observation for 1-step ahead perturbation plus, 0.24 seconds per observation for 2-step ahead perturbation plus, and 11 seconds per observation for 3-step ahead perturbation plus. The processor used is an Intel (R) Core(TM) i7-2600 CPU @ 3.40 GHz. Computing times for the alternatives are much smaller. For example, second-order pruned perturbation takes less than 0.0001 seconds per observation.

not so accurate first-order perturbation approximation to close the system.

5.2 Accuracy of approximations to the matching model

As in section 3.3, we consider two values for the wage rate, namely $w = 0.96$ and $w = 0.973$. At the lower value of w the regular second-order perturbation approximation is stable and at the higher value of w it is not. Accuracy measures and summary statistics for the different numerical approximations are given in tables 1 and 2, respectively. The following observations can be made.

First, the error measures are much smaller than those reported for the previous model. The reason is that the variance of the driving process is much smaller. By construction, the model generates plausible employment volatility (relative to labor productivity) when $w = 0.973$. In particular, the employment rate fluctuates between 89.8% and 96.4% according to the "true" solution. The model generates too little volatility in the employment rate when $w = 0.96$.

Second, first-order perturbation again performs very poorly. When $w = 0.973$, the first-order perturbation approximation predicts that the mean employment rate is equal to 94.9% while in fact it should be 93.1%. Making on average a mistake of 1.8 percentage point when the range of generated values is only 6.6 percentage points is very troubling.

Third, when $w = 0.973$ the regular second-order perturbation approximation explodes and is obviously outperformed by pruned perturbation. When $w = 0.96$ and the regular second-order *just* does not explode, it outperforms pruned perturbation by far. In fact, it performs very well.

Fourth, the results for second-order pruned perturbation are substantially better than those for first-order perturbation. Nevertheless, the time paths generated by second-order pruned perturbation are inaccurate in some important dimensions. In particular, it underestimates the depth of recessions. Consider the case when $w = 0.973$. The worst that can happen according to the second-order pruned perturbation approximation is a drop in the employment rate to 91.4%, while the actual minimum is 89.8%. In contrast, second-order pruned perturbation predicts the maximum value quite well. It predicts a peak of

96.35% whereas our accurate projection method predicts a peak of 96.44%. Consistent with these observations, second-order pruned perturbation underestimates the standard deviation of employment with 35%.

Fifth, the perturbation-plus procedure does very poorly unless the number of forward looking steps is very high. For the Brock-Mirman model, we found that even the one-step ahead perturbation-plus procedure delivered a substantial improvement over first-order perturbation. For this model, the one-step ahead perturbation-plus procedure leads to only a small improvement. As documented in table 1, even the nine-step ahead perturbation-plus procedure performs worse than the second-order pruned perturbation approximations. To get performance that is comparable to that of the second-order pruned perturbation approximation, the number of forward looking steps has to be at least 15.^{51,52}

The perturbation-plus procedure does much better if we do not use the first-order perturbation approximation to close the system, but use instead a policy function that is linear in n_{-1} but for which the coefficients depend on z . That is, for each of the two values of z we use a separate linear policy function. In fact, even the one-step ahead perturbation-plus procedure outperforms the second-order pruned perturbation by far. The two linear functions were obtained with a projection method. It is not difficult to obtain these two linear functions and one always achieves a stunning improvement as long as one conditions on the value of z . But we do not want to pursue this modification, since the idea of the perturbation-plus procedure is that one does not have to worry about setting up a grid, which is part of the projection method procedure.

⁵¹Even with the simplification discussed in online appendix B.2, it takes 102 seconds per observation to run 15-step ahead perturbation plus. The processor used is an Intel (R) Core(TM) i7-2600 CPU @ 3.40 GHz.

⁵²We also considered some cases when z had continuous support. Interestingly, one-step ahead perturbation plus then substantially improves upon first-order perturbation and also does much better than the two second-order perturbation approximations. The advantage of discrete support for second-order perturbation approximations is that it limits the maximum deviations from the steady state, which is useful given that the nonlinear dependence of n on z is not captured well with a second-order perturbation approximation.

5.3 Accuracy of approximations to the modified Deaton model

As documented in table 1, there are substantial differences between the time paths generated by the first-order perturbation approximation and the time path generated by the accurate projection method. For example, according to the accurate projection approximation, asset holdings range between -0.094 and 0.848 , and mean asset holdings are equal to 0.085 . In contrast, according to the first-order perturbation approximation, asset holdings range between -0.184 and 0.344 and the mean is equal to 0.029 . The first-order perturbation approximation has a maximum error of 137% and an average error of 44%.

Improvements are obtained by the nonlinear approximations, except by the regular second-order perturbation approximation, because it generates a time path that is not stable. The maximum error of the second-order pruned perturbation approximation is equal to 127%, only a slight improvement over the one obtained by the first-order perturbation approximation. But the average error is substantially less, namely 12.2%. The average error of the two-step ahead perturbation-plus procedure is slightly smaller than the one for second-order pruning, but the maximum error is substantially smaller, namely 64.6%.

Table 2 documents that there are also substantial differences in the properties of the generated time paths. The two-step ahead perturbation-plus procedure predicts a standard deviation that is 10% below the true value and second-order pruning underpredicts the standard deviation by 16%. The approximations have most difficulty in following the true time path in extreme situations. Interestingly, the perturbation-plus procedure has problems with the peaks and pruning has problems with the troughs.

Panel A of Figure 9 plots that part of the generated time path where the perturbation-plus procedure makes the biggest error and Panel B plots that part where pruning makes the biggest error. First consider Panel A. The figure makes clear that the perturbation-plus procedure is not as bad as the results reported above indicate. At the peak, the two-step ahead perturbation-plus approximation underestimates the maximum achieved level of asset holdings with 16%. Given that the true maximum is ten times as big as the mean asset holdings, this error is not that worrisome. In particular, the perturbation-

plus approximation also predicts an enormous increase and some time before and some time after this unusual situation the perturbation-plus procedure does track the accurate time path reasonably well. The maximum error for the two-step ahead perturbation-plus approximation is equal to 64%. Interestingly, the maximum error is not obtained at the peak but just after the peak, when observations have fallen back to normal levels.

Now turn to Panel B of Figure 9 that plots that part of the time path during which the second-order pruning procedure makes the biggest error. This occurs during a serious downturn. Similar to the case of panel A, the pruned perturbation approximation follows the accurate time path quite closely before and after this extreme period. The error made by second-order pruning may be a bit more serious than the error made by the perturbation-plus approximation in the top panel. First, the trough highlighted in Panel B is much closer to the mean level of asset holdings than the peak in Panel A. Moreover, the generated pattern is actually a bit different as well. After the time path generated with pruned perturbation has reached its minimum value in period 4362, it starts an upward movement with only a minor downward blip, whereas the accurate time path reaches a considerably lower minimum several periods later.

6 Concluding comments

In this paper, we have focused on nontrivial numerical problems. In particular, we have looked at models for which regular second-order perturbation generates time paths that reach that part of the state space where the derivative of the second-order approximation has the wrong sign and we have looked at models where the parameter values are outside the set that ensure stability.

The first-order perturbation approximation avoids both problems, but these are models for which nonlinearities matter. The two perturbation-based alternatives to regular perturbation approximations generate stable time paths and incorporate nonlinearities. These methods can deliver substantial improvements over linear approximations. In the examples considered in this paper, the proposed alternatives provide a good *qualitative* insight in how the true time path differs from the path generated by first-order pertur-

bation. In most cases, however, there are important quantitative differences between the time paths generated by the approximations and the accurate solution method. The two exceptions are the three-step ahead perturbation-plus procedure when applied to the Brock-Mirman (with high variance) and the modified Deaton model.

Thus, if the user is interested in precise quantitative properties of the model, then it is important to evaluate the quality of the approximation with an accuracy test. It is always a good idea to do this, but this is especially the case when solving models that are such that regular perturbation approximations generate unstable time paths.

To conclude, we would like to point out that there is nothing in the perturbation plus procedure that requires future behavior to be described by the first-order perturbation approximation. We used this approximation as input, because we were mainly interested to see whether small modifications of the (simple) first-order perturbation solution could lead to substantial improvements. But our procedure can be applied to any numerical approximation, independent of the method used to obtain it.

The one-step ahead procedure is related to the "dynamic Euler equation" accuracy test used in Den Haan (2010). The idea is the following. Data generated by our one-period ahead modification are more accurate than data generated *directly* by the input approximation, which our one-period ahead modification only uses *indirectly* via next period's behavior. This gap between the two data series provides an accuracy measure. If the approximation is accurate, then the two data series will be close to each other. That is, if a numerical approximation is accurate, then applying our one-step ahead modification should result in a time series that is virtually identical to the time series simply generated by the numerical approximation itself. Note that this would happen if the true rational expectations solution would be used to describe next period's behavior.

Acknowledgement

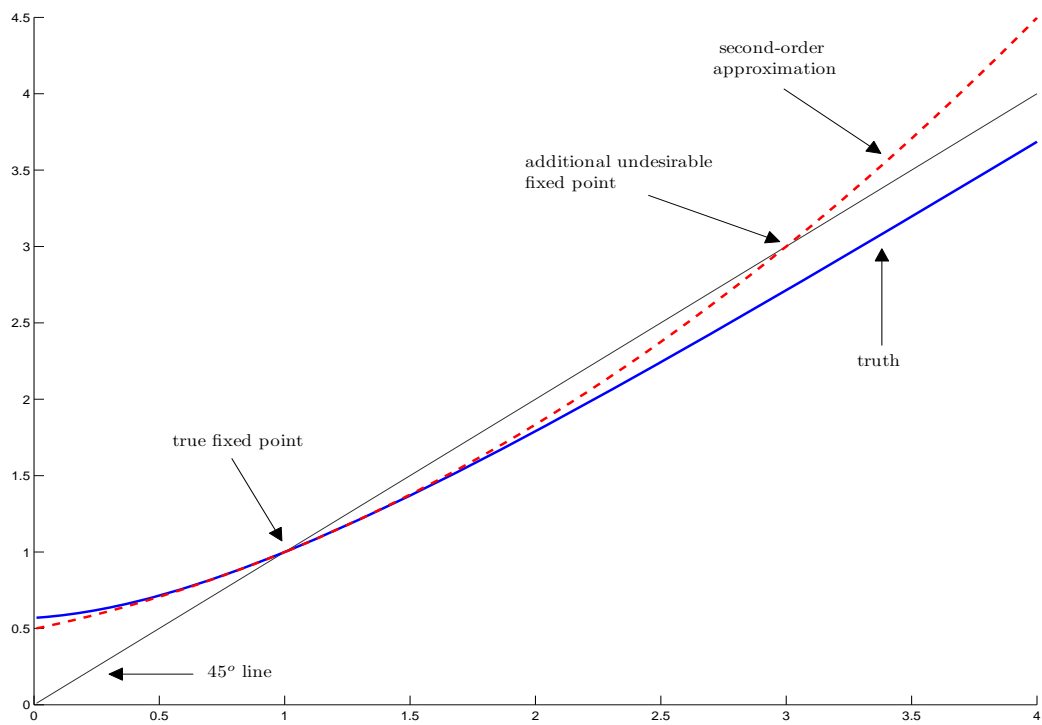
We would like to thank Fabio Canova, Ken Judd, Michel Juillard, Ondra Kamenik, Giovanni Lombardo, Tarik Oçaktan, the editor Paul Klein, and two anonymous referees for useful comments. Den Haan thanks the Netherlands Organisation for Scientific Research (NWO) for financial support.

References

- ANDREASEN, M. M. (2012): “Non-Linear DSGE Models, The Central Difference Kalman Filter,” *Journal of Applied Econometrics*, forthcoming.
- COLLARD, F., AND M. JUILLARD (2001): “Accuracy of Stochastic Perturbation Methods: The Case of Asset Pricing Models,” *Journal of Economic Dynamics and Control*, 25, 979–999.
- DE WIND, J. (2008): “Punishment Functions,” Unpublished manuscript, University of Amsterdam.
- DEATON, A. (1991): “Saving and Liquidity Constraints,” *Econometrica*, 59, 1221–1248.
- DEN HAAN, W. J. (2010): “Comparison of Solutions to the Incomplete Markets Model with Aggregate Uncertainty,” *Journal of Economic Dynamics and Control*, 34, 4–27.
- DEN HAAN, W. J., AND J. DE WIND (2009): “How Well-Behaved are Higher-Order Perturbation Solutions?,” DNB Working Paper No. 240.
- DEN HAAN, W. J., AND G. KALTENBRUNNER (2009): “Anticipated Growth and Business Cycles in Matching Models,” *Journal of Monetary Economics*, 56, 309–327.
- DOH, T. (2011): “Yield Curve in an Estimated Nonlinear Macro Model,” *Journal of Economic Dynamics and Control*, 35, 1229–1244.
- FAHR, S., AND F. SMETS (2010): “Downward Wage Rigidities and Optimal Monetary Policy in a Monetary Union,” *Scandinavian Journal Of Economics*, 112, 812–840.
- FAIR, R. C., AND J. B. TAYLOR (1983): “Solution and Maximum Likelihood Estimation of Dynamic Nonlinear Rational Expectations Models,” *Econometrica*, 51, 1169–1185.
- FERNÁNDEZ-VILLAYERDE, J., P. GUERRÓN-QUINTANA, J. F. RUBIO-RAMÍREZ, AND M. URIBE (2011): “Risk Matters: The Real Effects of Volatility Shock,” *American Economic Review*, 101, 2530–2561.

- GAGNON, J. E. (1990): “Solving the Stochastic Growth Model by Deterministic Extended Path,” *Journal of Business and Economic Statistics*, 8, 35–36.
- HAGEDORN, M., AND I. MANOVSKII (2008): “The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited,” *American Economic Review*, 98, 1692–1706.
- HALL, R. E. (2005): “Employment Fluctuations with Equilibrium Wage Stickiness,” *American Economic Review*, 72, 50–65.
- JUDD, K. L. (1998): *Numerical Methods in Economics*. The MIT Press, Cambridge, Massachusetts.
- KIM, J., S. KIM, E. SCHAUMBURG, AND C. A. SIMS (2008): “Calculating and Using Second-Order Accurate Solutions of Discrete Time Dynamic Equilibrium Models,” *Journal of Economic Dynamics and Control*, 32, 3397–3414.
- KRUSELL, P., AND A. A. SMITH, JR. (1998): “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 106, 867–896.
- LOMBARDO, G. (2010): “Approximating DSGE Models by Series Expansions,” Unpublished manuscript, European Central Bank.
- PETRONGOLO, B., AND C. PISSARIDES (2001): “Looking Into the Black Box: A Survey of the Matching Function,” *Journal of Economic Literature*, 39, 390–431.


Figure 1: Perturbation approximations and instability



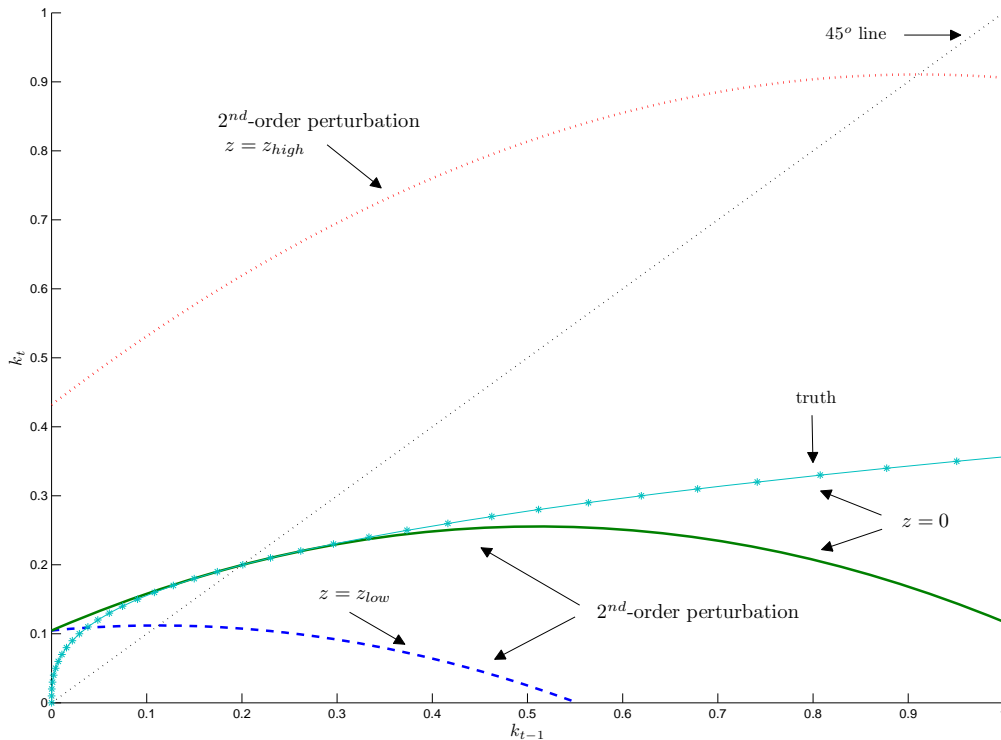
Notes: This figure plots the function $f(x_{-1})$ described in section 1 and its second-order Taylor-series approximation.

Figure 2: 2nd-order perturbation approximations for Brock-Mirman model

A: Approximation for $k_t = k(k_{t-1}, z_t)$; $\sigma_z = 0.007$




B: Approximation for $k_t = k(k_{t-1}, z_t)$; $\sigma_z = 0.2$



Notes: This figure plots the 2nd-order perturbation approximations for the Brock-Mirman model in levels (not in logarithms) for different values of z .


Figure 3: A 2nd-order perturbation approximation for neoclassical growth model




Notes: Capital is solved from the true budget constraint using the perturbation approximation for consumption. When $\gamma = 1$, then the model is the Brock-Mirman model.

Figure 4: "Truth" and 2nd-order perturbation approximation for matching model

A: $w = 0.96$.




B: $w = 0.973$.




Notes: The point of this figure is to show that the second-order perturbation solution when $z = -\zeta$ does not have a fixed point unless w is sufficiently low (as in panel A).

Figure 5: 2nd-order perturbation approximations for the Deaton model




Notes: This figure shows that the second-order perturbation approximation for the policy function of the Deaton model is not globally stable. The graph plots $E[x_{t+1}|x_t]$ as a function of x_t , where x_t is cash on hand (the only state variable of the model). Cash on hand is equal to lagged asset holdings plus current-period income. The bottom panel gives the distribution for x_t according to the "truth".

Figure 6: Higher-order perturbation and the modified Deaton model



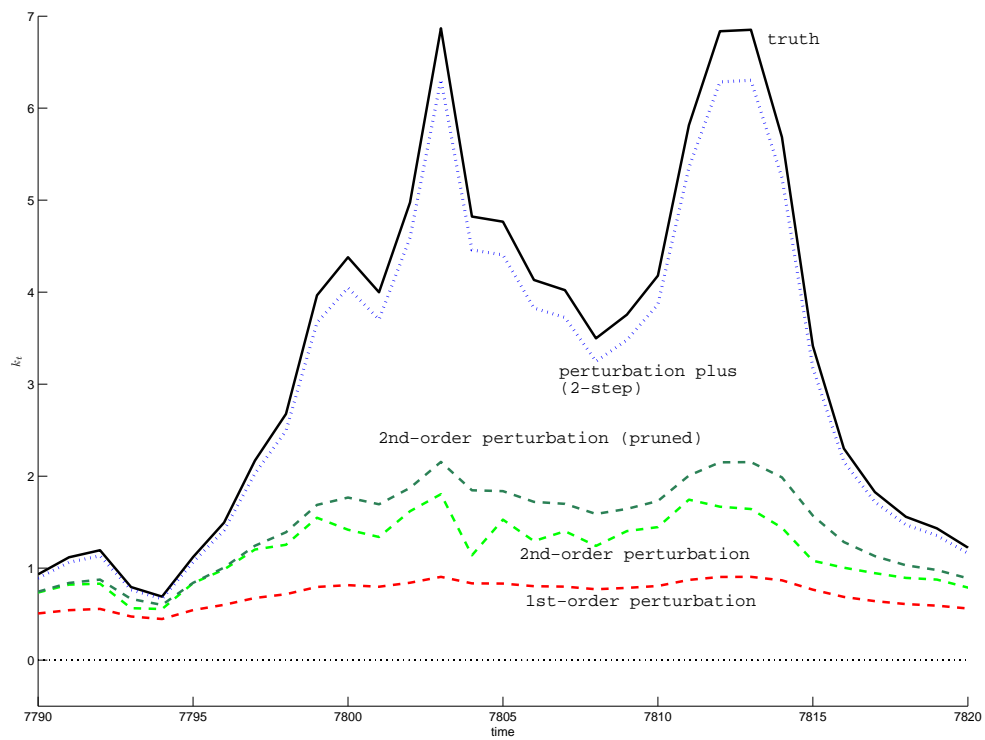
Notes: The graph plots $E[x_{t+1}|x_t]$ as a function of x_t , where x_t is cash on hand (the only state variable of the model). Cash on hand is equal to lagged asset holdings plus current-period income. The bottom panel gives the distribution for x_t according to the "truth".

Figure 7: Data simulated with 2nd-order pruned perturbation approximation;
 $k_t^{(2)} - k_{t-1}^{(2)}$ as a "function" of $k_{t-1}^{(2)}$



Notes: $k_t^{(2)}$ is the capital stock generated by the pruning procedure. $k_t^{(2)}$ is not a function of $k_{t-1}^{(2)}$ and z_t , since the pruning procedure introduces additional state variables.

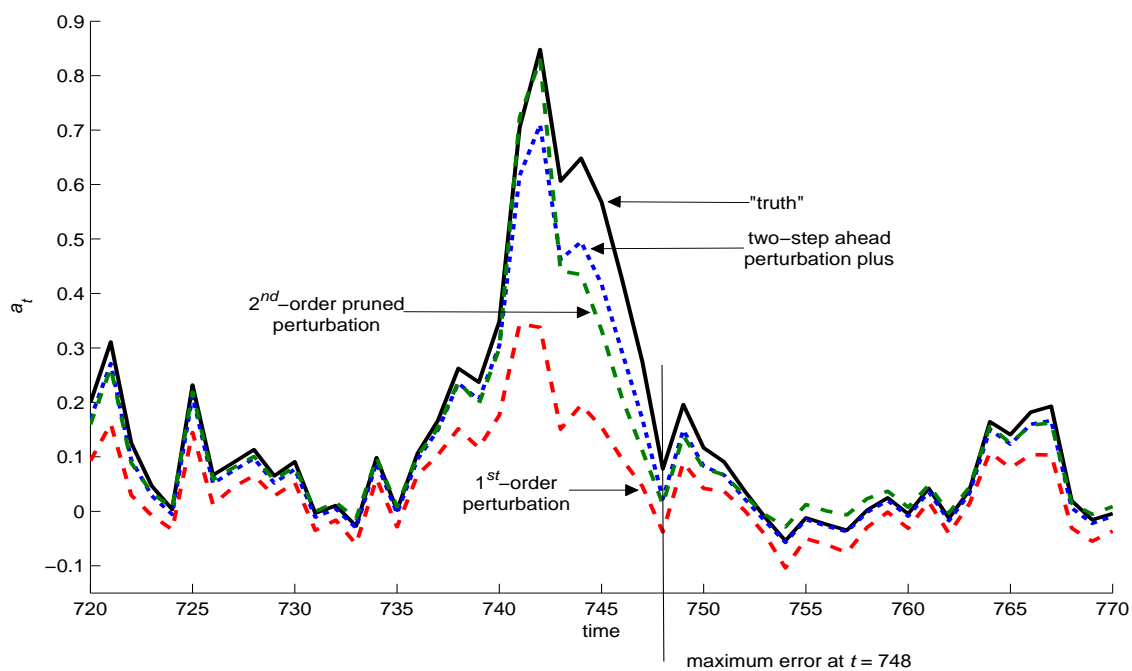
Figure 8: Capital stock time paths for the Brock-Mirman model



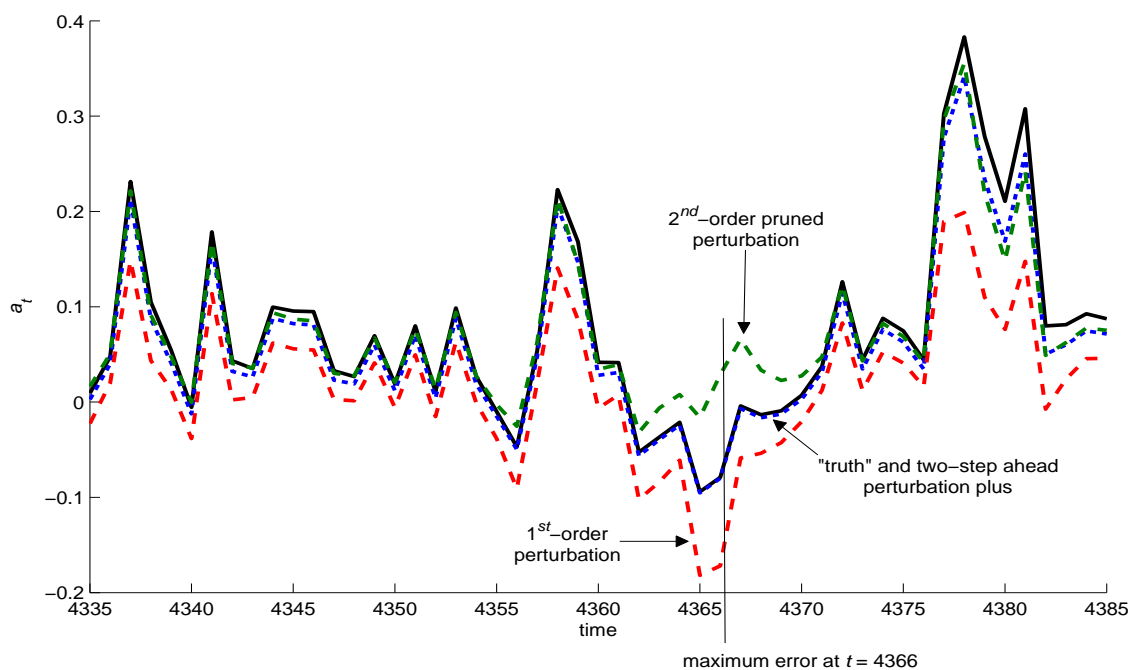
Notes: This graphs plots the different approximations in that part of the generated time series where the 2-step ahead perturbation-plus approximation attains its largest error.

Figure 9: Asset holding time paths for the modified Deaton model

A: When two-step ahead perturbation plus attains largest error



B: When 2nd-order pruned perturbation attains largest error



Notes: Panel A (B) plots that part of the generate time series where the two-step ahead perturbation-plus approximation (2nd-order pruned perturbation) attains its largest error.

Table 1: Differences (in %) between approximations and "truth"

Brock-Mirman (capital stock)				
	$\sigma_z = 0.1$		$\sigma_z = 0.2$	
	max	mean	max	mean
perturbation + 1	16.7	2.7	25.3	5.4
perturbation + 2	5.2	1.0	8.1	1.9
1 st -order pert.	76.1	8.0	142.8	19.2
2 nd -order pert.	31.0	1.9	76.4	7.9
2 nd -order pruning	47.9	2.0	193.8	8.8
matching (employment)				
	$w = 0.96$		$w = 0.97$	
	max	mean	max	mean
perturbation + 1	0.97	0.67	3.09	1.82
perturbation + 1*	0.30	0.12	0.44	0.13
perturbation + 9**	0.66	0.45	2.38	1.38
1 st -order pert.	1.02	0.69	3.20	1.87
2 nd -order pert.	0.06	0.02	∞	∞
2 nd -order pruning	0.36	0.22	1.79	0.94
modified Deaton (asset holdings)				
	max	mean		
perturbation + 1	97.9	22.1		
perturbation + 2	64.6	11.1		
1 st -order pert.	136.9	44.4		
2 nd -order pert.	∞	∞		
2 nd -order pruning	127.7	12.2		

Notes: "truth" refers to the true policy rule for the Brock-Mirman model and to a very accurate projection method for the other two models. The * indicates that two separate linear rules are used for the two values of z and ** indicates that the approximation is calculated as described in appendix B.1. Results are based on a time path of 10,000 observations.

Table 2: Model properties according to approximations and "truth"

Brock-Mirman (capital stock)								
	$\sigma_z = 0.1$				$\sigma_z = 0.2$			
	E[k]	σ_k	min	max	E[k]	σ_k	min	max
truth	0.227	0.123	0.037	1.170	0.334	0.444	0.007	6.868
perturbation + 1	0.218	0.114	0.010	0.983	0.296	0.360	0.000	5.158
perturbation + 2	0.224	0.120	0.029	1.110	0.322	0.417	0.004	6.312
1 st -order pert.	0.200	0.099	-0.136	0.552	0.201	0.198	-0.471	0.905
2 nd -order pert.	0.223	0.107	0.089	0.838	0.283	0.232	0.089	1.804
2 nd -order pruning	0.225	0.108	0.010	0.865	0.299	0.253	0.100	2.155
matching (employment)								
	$w = 0.96$				$w = 0.97$			
	E[n]	σ_n	min	max	E[n]	σ_n	min	max
"truth"	0.943	0.0175	0.925	0.961	0.931	0.0315	0.898	0.964
perturbation + 1	0.949	0.0151	0.934	0.965	0.948	0.0218	0.927	0.972
perturbation + 1*	0.941	0.0178	0.924	0.961	0.932	0.0318	0.899	0.966
perturbation + 9**	0.947	0.0157	0.931	0.964	0.944	0.0236	0.920	0.969
1 st -order pert.	0.949	0.0151	0.934	0.966	0.949	0.0217	0.928	0.972
2 nd -order pert.	0.942	0.0176	0.925	0.961	∞	∞	∞	∞
2 nd -order pruning	0.944	0.0155	0.928	0.961	0.937	0.0228	0.915	0.964
modified Deaton (asset holdings)								
	E[a]	σ_a	min	max				
"truth"	0.085	0.098	-0.094	0.848				
perturbation + 1	0.057	0.080	-0.101	0.583				
perturbation + 2	0.071	0.088	-0.097	0.710				
1 st -order pert.	0.029	0.070	-0.184	0.344				
2 nd -order pert.	∞	∞	-	∞				
2 nd -order pruning	0.079	0.082	-0.043	0.829				

Notes: "truth" refers to the true policy rule for the Brock-Mirman model and to a very accurate projection method for the other two models. The * indicates that two separate linear rules are used for the two values of z and ** indicates that the approximation is calculated as described in appendix B.1. Results are based on a time path of 10,000 observations.