E C O N O M I C S B U L L E T I N

On the accuracy of low-order projection methods

Paul Pichler
University of Vienna

Abstract

We use low-order projection methods to compute numerical solutions of the basic neoclassical stochastic growth model. We assess the quality of the obtained solutions, and compare them to numerical approximations derived with first and second-order perturbation techniques. We show that projection methods perform surprisingly poor when the degree of approximation is very low, and we provide some intuition behind this finding.

A previous, longer version of this paper has been circulated under the title "Evaluating Approximate Equilibria of Dynamic Economic Models". I thank Gerhard Sorger and Peter Hackl for many helpful comments and suggestions. Financial support from the Austrian Science Fund (FWF) under project numbers P17886 and P19686 is gratefully acknowledged.
Citation: Pichler, Paul, (2007) "On the accuracy of low-order projection methods." *Economics Bulletin*, Vol. 3, No. 50 pp. 1-8
Submitted: May 2, 2007. **Accepted:** October 10, 2007.
URL: http://economicsbulletin.vanderbilt.edu/2007/volume3/EB-07C60003A.pdf

1. Introduction

Projection methods are becoming an increasingly popular approach for solving dynamic economic models.[1] The main reason behind this development is the extraordinary robustness and flexibility they offer. In particular, the projection approach works well in many applications where other methods typically fail, as for example, to solve models with occasionally binding inequality constraints.[2] Furthermore, it is well documented in recent papers by Aruoba, Fernandez-Villaverde, and Rubio-Ramirez (2006) and Heer and Maussner (2004) that, when implemented with high-degree polynomials, projection methods allow to derive very precise approximations to the equilibria of dynamic economic models.

In many of the recent applications, however, the complexity of the considered models places a constraint on the feasible order of approximation within the projection algorithm. Precisely, in models with many state variables, high-order accurate solutions are computationally too demanding. This is due to the well-known *curse of dimensionality*. Researchers are thus left with using low-order implementations of the projection approach. Still, very little is known about the quality of these low-order projection solutions. It is the aim of this note to provide some evidence on this issue.

We proceed as follows. Section 2 briefly discusses the model we use to judge the quality of approximate solutions. Section 3 gives a short illustration of the numerical solution methods we employ. Section 4 presents the results of our evaluation exercise. Finally, we provide some conclusions in section 5.

2. The model

We evaluate the quality of low-order projection solutions using the basic neoclassical stochastic growth model as a testing ground. We basically choose this model for two reasons. First, several solution methods have already been tested on this model, such that our results can easily be compared to other studies.[3] Secondly, the model forms the core of many models used in modern macroeconomics, such that lessons learnt with this simple framework are likely to hold also (to some extent) in other setups.

As the model is very well-known, we restrict ourselves to presenting the main formulas and the competitive equilibrium. The model economy is populated by a single representative household. Time is discrete, i.e. $t \in \{0, 1, 2, \ldots\}$. The household chooses sequences of consumption $\{c_t\}_{t=0}^{\infty}$, investment $\{i_t\}_{t=0}^{\infty}$, output $\{y_t\}_{t=0}^{\infty}$ and capital $\{k_{t+1}\}_{t=0}^{\infty}$ such that it maximizes the expected discounted utility derived from its current and future consumption stream. The maximization is subject to constraints posed by a production function, a budget constraint, a law of motion for the exogenous productivity level, and by non-negativity requirements for consumption and capital. We assume a CES utility function and a Cobb-Douglas production function. Formally, the household solves

$$\max_{c_t,\, i_t,\, y_t,\, k_{t+1}} \quad E_0 \sum_{t=0}^{\infty} \beta^t \frac{1}{1-\gamma} c_t^{1-\gamma} \tag{1}$$

[1] For recent applications see Bohacek and Kejak (2005), Ortigueira (2006), Siu (2004), Reiter (2005), Gapen and Cosimano (2005), Jones, Manuelli, and Siu (2005) and Dueker, Fischer, and Dittmar (2004).

[2] See Judd (1992) and Christiano and Fisher (2000).

[3] See, for example, Taylor and Uhlig (1990).

subject to

$$\begin{aligned} y_t &= a_t k_t^{\alpha} \\ i_t &= y_t - c_t \\ k_{t+1} &= (1-\delta)k_t + i_t \\ \log a_{t+1} &= \rho \log a_t + \sigma_\varepsilon \varepsilon_{t+1} \\ \varepsilon_{t+1} &\sim N(0,1) \\ c_t, k_t &\geq 0 \quad \text{for all } t \end{aligned}$$

The parameters and variables have the following interpretation: a_t denotes an exogenous technology level, $\beta \in (0,1)$ is the household's time discount factor, $\alpha \in (0,1)$ describes capital's share in output, $\gamma > 0$ denotes the household's risk aversion, $\delta \in (0,1)$ is the depreciation rate of capital, ϵ_{t+1} denotes a shock to technology at time $t+1$, the parameter σ_ε captures the standard deviation of technology shocks, and ρ governs their persistence.
It is well known that the necessary condition for optimality in this model is given by the Euler equation

$$c_t^{-\gamma} = \beta E_t c_{t+1}^{-\gamma} [\alpha a_{t+1} k_{t+1}^{\alpha-1} + (1-\delta)]. \tag{2}$$

Above, E_t denotes the expectation operator conditional upon the information available to the decision maker at time t. When augmented with a transversality condition

$$\lim_{t\to\infty} \beta^t c_t^{-\gamma} k_{t+1} = 0, \tag{3}$$

the Euler equation (2) becomes sufficient for optimality.[4]
The full system of necessary equilibrium conditions can be summarized by an operator

$$\mathcal{R}(\cdot) = \begin{pmatrix} E_t \beta {c_{t+1}}^{-\gamma} [\alpha a_{t+1} {k_{t+1}}^{\alpha-1} + (1-\delta)] - c_t^{-\gamma} \\ y_t - a_t k_t^{\alpha} \\ y_t - c_t - i_t \\ (1-\delta)k_t + i_t - k_{t+1} \\ \rho \log a_t - \log a_{t+1} \end{pmatrix} = 0.$$

The solution of our model is the zero of the operator $\mathcal{R} : B_1 \to B_2$, where B_1 and B_2 are function spaces. We denote this solution by a vector-valued function $f : D \subset \mathbb{R}^2_+ \to \mathbb{R}^5$. The function f maps the state space (spanned by the state vector $x_t = (k_t, a_t)'$) into decision outcomes for consumption, investment and output, as well as the next period capital stock and technology level.

3. Approximate equilibria

Since we cannot compute f analytically, we use numerical methods to obtain approximate solutions. An approximate solution $\hat{f}$ is characterized by the requirement $\mathcal{R}(\hat{f}) \approx 0$. Alternatively,

[4] See, for example, Stokey and Lucas (1989) for a proof of the sufficiency

we can define an approximate equilibrium as the solution of $\hat{\mathcal{R}}(\hat{f}) = 0$, where we have replaced the operator $\mathcal{R}$ by a (computable) operator $\hat{\mathcal{R}}$. Note that $\hat{\mathcal{R}}$ implicitly defines the norm we use to judge the size of $\mathcal{R}(\hat{f})$.

When using projection methods, we have several degrees of freedom in the choice of $\hat{f}$ and $\hat{\mathcal{R}}$. A general kth order projection approach uses orthogonal polynomials up to order k to construct $\hat{f}(x;\kappa)$, where x is the vector of state variables and κ summarizes the unknown coefficients in the approximate decision rules. Note that typically only a few elements of f have to be approximated by parametric forms, since the remaining elements can be computed directly from the operator $\mathcal{R}$. In our example, we parameterize only the decision rule for consumption, whereby we use a complete basis of Chebyshev polynomials. Conditional on $\hat{c}(x;\kappa)$, we derive the remaining elements of $\hat{f}$ directly from the last four conditions in $\mathcal{R}(\hat{f}) = 0$.

Having constructed a functional form $\hat{f}(x;\kappa)$, the problem of solving $\hat{\mathcal{R}}(\hat{f}) = 0$ boils down to finding an appropriate parameter vector κ. Projection methods achieve this by solving

$$\int_D \mathcal{R}(x;\hat{f},\kappa)\omega_i(x)dx = 0 \quad i = 1,\ldots,\#\kappa \tag{4}$$

The functions $\omega_i(x)$ are weighting functions that define the norm used to judge the size of $\mathcal{R}(x;\hat{f},\kappa)$. The three most popular choices correspond to *Least squares*, *Chebyshev Collocation*, and *Galerkin* methods. In this paper, we follow a Galerkin approach.[5]

In section 4, we compare the quality of approximate solutions obtained via low-order Galerkin projection methods with solutions obtained with the popular perturbation approach. The latter replaces $\mathcal{R}$ with its kth order Taylor series expansion around the non-stochastic steady state, and inverts $\hat{\mathcal{R}}(\hat{f}) = 0$ to obtain kth-order accurate decision rules for all variables. The main strength of this approach lies in its speed. Even for large-scale models, perturbation solutions can typically be derived within a couple of seconds. Its arguably main weakness is that the perturbation approach is not universally applicable. For example, it cannot be applied to solve models with occasionally binding inequality constraints.

4. Results

The main purpose of this paper is to investigate the quality of low-order projection solutions. To this end, we solve our model with Galerkin methods of order 1, 2, and 4, and we compare these solutions with approximations obtained with first and second-order perturbation methods. Overall, we thus evaluate 5 different numerical procedures.

Since the model's structural parameters strongly affect the non-linearity of the problem and the deviations of the model's variables around the steady state, they are likely to have a non-negligible effect on our results. To guarantee robustness, we thus consider 500 different calibration sets for our evaluation. We derive these parameters by drawing randomly from postulated uniform distributions, which are given in table 1. These distributions are chosen to cover the parameter range typically considered in the literature.

For each parameter constellation, we solve the model using any of our five approaches, and we construct artificial time series for all model variables. We use these series to compute

[5]Further details are provided in the working paper version, Pichler (2005).

Table 1: Postulated Distributions for Structural Parameters

α	β	γ	δ	ρ	σ_ε
$U[0.24, 0.36]$	$U[0.95, 0.99]$	$U[1.5, 3]$	$U[0.025, 0.1]$	$U[0.85, 0.95]$	$U[0.01, 0.1]$

Euler equation errors as suggested by Judd (1992). The average absolute and the maximum absolute error across our simulations will serve as the main indicators of numerical accuracy. Furthermore, we compare the averages and standard deviations of the simulated consumption and capital series.

Let us first discuss the Euler equation measures. We summarize the base 10 logarithm of the average error ($\log_{10} ||E||_1$) and the maximum error ($\log_{10} ||E||_\infty$) associated with a specific approach in table 2. Furthermore, we list the average computing time (in seconds) required by every method. We observe that for the case of a linear approximation, Galerkin projection performs only poorly. The associated mean absolute euler equation error is approximately seven times larger as compared to first-order perturbation. In particular, it states that on average the household makes a mistake of 0.4% in its consumption decision. This error is by far larger than what is typically accepted by a researcher.

Table 2: Mean Computer Time and Euler Equation Errors

	$\log_{10} ||E||_1$	$\log_{10} ||E||_\infty$	CPU
Galerkin 1	-2.3803	-1.9831	1.4121
Galerkin 2	-4.3595	-3.7939	3.2370
Galerkin 4	-6.1182	-5.7091	8.1613
Perturbation 1	-3.2248	-2.2989	0.3587
Perturbation 2	-4.8259	-3.6626	0.4362

In the case of a second-order approximation, perturbation methods still outperform Galerkin, although the difference is not as big as in the case of linear decision rules. A fourth-order implementation of the projection approach, on the other hand, delivers very accurate results. The resulting Euler equation errors are of several magnitudes smaller than those of a second-order perturbation solution. Unfortunately, the computing time necessary to solve our model with a fourth-order accurate projection method is substantially higher than with second-order perturbation. In many applications, a fourth order implementation may not be feasible.

In the following, we investigate whether the dynamic properties of the model are affected by the choice of solution method. In particular, we analyze whether low-order approximations distort the first and second moments of our simulated consumption and capital series. The results in table 3 provide an affirmative answer to this question. The time series simulated from the model solved by a linear Galerkin method are substantially distorted. The average capital stock is approximately 20% smaller than implied by higher order approximations. Similarly, we observe that the mean capital stock associated with first-order perturbation is substantially biased, too (-5%). This results from the well known certainty equivalence property of linear perturbation

Table 3: First and Second Moments of Simulated Consumption and Capital Series

	Consumption		Capital	
	Mean	Std	Mean	Std
Galerkin 1	1.3697	0.2102	5.6038	1.4012
Galerkin 2	1.4151	0.2039	6.9427	1.7460
Galerkin 4	1.4155	0.2040	6.9510	1.7337
Perturbation 1	1.4003	0.2055	6.6450	1.6703
Perturbation 2	1.4177	0.2057	6.9861	1.7725

solutions: linearization around a non-stochastic steady state ignores precautionary saving, and thus, implies a smaller capital stock in our model. Finally, the results in table 3 show that all approximations of order higher than one deliver similar implications for the first and second moments of consumption and capital.

Taken together, tables 2 and 3 indicate a surprisingly poor performance of low-order projection methods. What is the explanation behind this result? Since, for a given order of approximation, projection and perturbation methods use the same functional forms for the approximate decision rules, the reason for the better performance of the perturbation approach must lie in the norm used to determine the unknown coefficients. Indeed, both approaches differ crucially in this aspect. Perturbation methods choose the coefficients such that the approximation becomes exact in the non-stochastic steady state, and thus, they deliver approximations which are very accurate in the neighborhood of this point. Projection methods, on the other hand, use global criteria. In particular, they minimize a weighted approximation error over the entire state space, whereby they put a heavy weight on the outer regions. If the functional form for the decision rules is a low-order polynomial, this implies that the approximation is relatively poor, as compared to perturbation, in the neighborhood of the steady state. Our results demonstrate that, only if higher-order polynomials are in use, projection methods achieve a good fit in the neighborhood of the steady state. Since our model stays relatively close to the steady state during the simulation, it is not surprising that perturbation delivers better results for low-order approximations.

Let us finally emphasize this argument using the following simple example. Assume that we want to obtain an approximate solution to the function f in the interval $[-1, 1]$, whereby f is defined by

$$f(x) - e^x + 1 = 0.$$

Obviously, the true solution is given by $f(x) = e^x - 1$. Figure 1 plots the errors $r(x) = \hat{f}(x) - f(x)$ for each of our considered numerical methods. We observe that the quality of perturbation approximations is very high in the neighborhood of $x^* = 0$, in particular, the approximations are better as compared to first and second-order projection methods. Only if we move far away from 0, $abs(x) > 0.5$, the low-order projection approximations are better. A projection solution of order 4, on the other hand, is able to generate an excellent approximation. The corresponding error is virtually zero across the entire state space.

Figure 1: *Approximation errors to the function* $f(x) = e^x - 1$

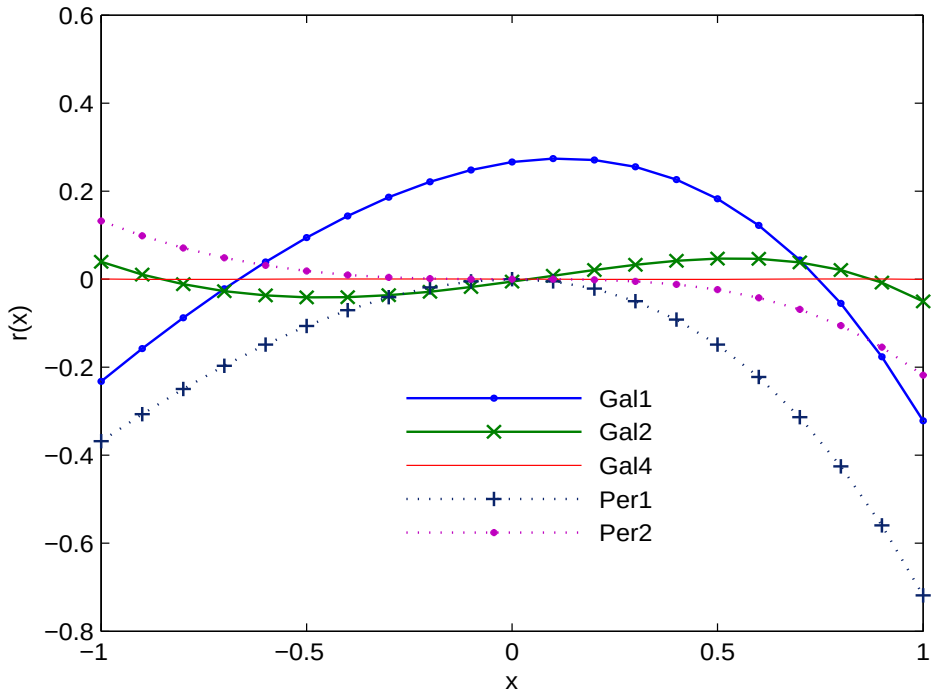

5. Concluding remarks

We have documented that, when implemented with a high degree of approximation, projection methods are an excellent way to solve dynamic economic models. If, however, the order of approximation is relatively low, then projection methods may perform only poorly. This will be true particularly for models where the economy stays relatively close to its steady state. A researcher following a low-order projection approach should have this result in mind, and thus she should always check the accuracy of the obtained numerical solution.

References

Aruoba, S. B., J. Fernandez-Villaverde, and J. Rubio-Ramirez (2006): "Comparing solution methods for dynamic equilibrium economies," *Journal of Economic Dynamics and Control*, 30(12), 2477–2508.

Bohacek, R., and M. Kejak (2005): "Projection Methods for Economies with Heterogeneous Agents," CERGE-EI Working Papers wp258.

Christiano, L. J., and J. D. M. Fisher (2000): "Algorithms for solving dynamic models with occasionally binding constraints," *Journal of Economic Dynamics and Control*, 24(8), 1179–1232.

Dueker, M. J., A. M. Fischer, and R. D. Dittmar (2004): "Stochastic capital depreciation and the comovement of hours and productivity," Discussion paper.

Gapen, M., and T. Cosimano (2005): "Solving Ramsey Problems with Nonlinear Projection Methods," *Studies in Nonlinear Dynamics & Econometrics*, 9(2), 1263–1263.

Heer, B., and A. Maussner (2004): "Projection methods and the curse of dimensionality," Working paper.

Jones, L. E., R. E. Manuelli, and H. E. Siu (2005): "Fluctuations in Convex Models of Endogenous Growth II: Business Cycle Properties," *Review of Economic Dynamics*, 8(4), 805–828.

Judd, K. L. (1992): "Projection methods for solving aggregate growth models," *Journal of Economic Theory*, 58(2), 410–452.

Ortigueira, S. (2006): "Markov-Perfect Optimal Taxation," *Review of Economic Dynamics*, 9(1), 153–178.

Pichler, P. (2005): "Evaluating Approximate Equilibria of Dynamic Economic Models," Vienna Economics Papers 0510, University of Vienna, Department of Economics.

Reiter, M. (2005): "Solving Models of Optimal Monetary and Fiscal Policy by Projection Methods," manuscript.

Schmitt-Grohe, S., and M. Uribe (2004): "Solving dynamic general equilibrium models using a second-order approximation to the policy function," *Journal of Economic Dynamics and Control*, 28(4), 755–775.

Siu, H. E. (2004): "Optimal fiscal and monetary policy with sticky prices," *Journal of Monetary Economics*, 51(3), 575–607.

Stokey, N. L., and R. E. Lucas (1989): *Recursive Methods in Economic Dynamics*. Harvard University Press, Cambridge, Massachusetts.

Taylor, J., and H. Uhlig (1990): "Solving Nonlinear Stochastic Growth Models: A Comparison of Alternative Solution Methods," *Journal of Business and Economic Statistics*, 8(1), 1–17.