Estimating the Accuracy of Numerical Solutions to Dynamic Optimization Problems*

Michael Reiter, Universitat Pompeu Fabra, Barcelona

April 2001

Keywords: dynamic optimization, Euler residuals, error bounds

Abstract

The paper considers stochastic dynamic optimization problems where a tentative numerical solution has been found. It uses the Euler residuals along simulated paths of the model to estimate the accuracy of the proposed solution. The main measure of accuracy is the the reduction in the criterion function from using the numerical rather than the exact solution. Estimates of the approximation error of the policy function can also be computed. The method can handle models with nonlinear inequality constraints. In a broad class of dynamic general equilibrium models, the method can be used to characterize the precision of a numerical solution by a bounded-rationality measure of the optimizing agents. The method is illustrated by an application to the one-dimensional stochastic growth model, where it is shown to provide quite precise estimates of the errors in the value and the policy function.

Address of the author:

Michael Reiter
Universitat Pompeu Fabra
Department of Economics and Business
Ramon Trias Fargas, 25–27
08005 Barcelona
Spain

e-mail: reiter@upf.es

*I am grateful to Ken Judd, Albert Marcet and participants of the Stanford SITE Workshop 2000 for many helpful suggestions. Special thanks to Jim Costain for detailed comments on earlier drafts of this paper. Support from the Spanish Ministry of Education under the grant PB98-1066 is gratefully acknowledged.

1 Introduction

With the necessity to solve more and more complicated dynamic optimization and dynamic general equilibrium models, economists usually resort to approximate numerical methods, including dynamic programming algorithms, linear-quadratic or higher order approximations, projection methods or even evolutionary algorithms. Very little is known about the general accuracy properties of these methods (except for dynamic programming, cf. Santos and Vigo-Aguiar, 1998), and it seems necessary to analyze the accuracy of a proposed solution in each specific case, to make sure that it is sufficient for the problem at hand. Often there are several numerical solution procedures available, and the task is to choose the best one.

Assume we have computed an approximate numerical solution to a given problem. If we deal with a dynamic optimization problem, we would like to know how far the solution is away from the exact solution in value terms. That means, what is the loss in the objective function that results from using the numerical approximation rather than the exact solution? The answer to this question would help us to decide whether it makes sense to invest more resources to obtain a more precise solution. If the problem is, more generally, an economic model where maximizing economic agents interact, with possible externalities, distortions etc., a relevant question is whether the proposed solution is an ϵ-equilibrium in the sense that every agent follows a boundedly rational strategy where her loss in utility is smaller than ϵ compared to the exactly optimal strategy, *given the dynamic behavior of other agents* (cf. Judd, 1992, Section 5.2, who defines the ϵ-equilibrium in terms of Euler residuals).

This paper develops a method to answer the above questions, without knowing the exact solution, and by solving a much simpler type of problem than the dynamic optimization problem under consideration. The method can be applied to most models used in macroeconomics. All that is needed is an estimate of the policy function of all agents, and the possibility to compute the Euler residuals from the model simulations using the estimated policy functions.

I think that the most essential information about the quality of a numerical solution is the value loss resulting from using that solution. In addition, one often wants to know additional aspects of accuracy, for example: does the numerical policy function approximate well the exact policy? does the solution describe correctly the reaction of policy functions to external parameters (uncertainty, taxes etc.)? The paper therefore also deals with estimating the error in the policy function.

Thorough comparisons of different numerical methods which are found in the literature (see, e.g., Taylor and Uhlig, 1990, Christiano and Fisher, 2000, Collard, Fève and

Juillard, 2000, Benitez-Silva, Hall, Hitsch, Pauletto and Rust, 2000) usually employ a variety of statistics and evaluation criteria. However, when the exact solution is not known (which is the relevant case), not many methods are available to estimate the accuracy (cf. Section 7 for a discussion of related approaches). The most useful instrument seems to be the Euler residual (see Judd, 1992, who also provides a way to normalize the residuals so that they have an economic interpretation). Santos (2000) derives theoretical relationships between the size of the Euler residuals and the deviations of the policy function and the value function from their exact values. The derived relationship is in the *supremum norm* of these variables.

The present paper follows Santos (2000) in measuring the approximation error in the value and the policy function, but it is complementary to that paper in the sense that it does not estimate the supremum (or maximum) of the error over the state space, but either the error at a particular point in the state space or the average error over a part of the state space. While Santos (2000) mainly derives theoretical bounds that are easy to compute but relatively loose, the present paper uses methods that are more computer intensive but give tight estimates. An important generalization of the present method compared to Santos (2000) is the fact that it can handle models with occasionally binding inequality constraints.

In numerical analysis it is more common to compute error bounds in supremum norms rather than average errors. If we can actually show that the supremum of the approximation error is below some required threshold level, we can have the highest level of confidence in a method. The disadvantage is that the supremum approach may often be non-operational if the accuracy of the solution at hand is not very high, since accuracy measures in supremum norms are very conservative. First, because the supremum estimates, necessarily based on conservative assumptions, usually give an estimated supremum error that is much bigger (by one or two orders of magnitude) than the true supremum error (cf. Section 2). This is a well known phenomenon in numerical analysis. Second, the supremum error criterion itself may be too demanding. Consider a numerical solution that provides the approximately right policy in almost all circumstances, but makes substantial mistakes under some circumstances. If these circumstances are very unlikely to ever occur, this numerical method may be easily good enough for a given purpose. In the same way, if we have to decide between two methods to use, we would often prefer the one with the lower average error over the one with the lower supremum error. Measures of average value loss, as presented in this paper, are therefore a useful alternative to supremum errors.

The plan of the paper is as follows. Section 2 provides a heuristic discussion of the relationship between Euler residuals and approximation errors. Section 3 describes the

class of optimization problems that we are going to study. Section 4 derives estimates of the approximation error in the value function, while the error in the policy function is treated in Section 5. Numerical examples are provided in Section 6. Section 7 concludes by discussing the conditions when it appears most appropriate to employ the method of this paper.

2 Euler residuals and approximation errors

The method of the present paper makes essential use of the Euler residuals. The purpose of this section is therefore to gain some intuitive understanding about the relationship between the Euler residuals and the value and policy error.

One potential problem is that the Euler equation alone is not a *sufficient* condition for a solution. An optimization problem usually has infinitely many paths that satisfy the Euler equation, i.e., have zero Euler residuals. Fortunately, it is well known that the Euler equation together with a transversality condition are sufficient for concave problems. A non-optimal solution which satisfies the Euler equation therefore violates the transversality condition or some feasibility condition. This implies that, at some point in time, it will do something that is either obviously non-optimal (in the standard consumption problem, letting the ratio of consumption to capital go to zero) or impossible (negative capital, growing at the rate of interest). Since a reasonable numerical method excludes this kind of anomalous behavior, the inaccuracy of the solution will *sooner or later* show up in the Euler residuals. We will detect the inaccuracy if we keep track of the Euler residuals over time.

Unfortunately, in order to get a precise estimate of the value and policy error, it is not enough to look at the maximum Euler residual. Let us illustrate this at a very simple case. Consider the deterministic finite horizon consumption problem with interest rate and discount factor equal to zero. The household maximizes

$$\sum_{t=1}^{T} u(c_t) \tag{1}$$

with $u'(c) > 0$ and $u''(c) < 0$, subject to the constraint

$$\sum_{t=1}^{T} c_t = K \tag{2}$$

with given initial capital K. The Euler residual for this problem is

$$u'(c_t) - u'(c_{t+1}) \tag{3}$$

and the optimal consumption path is

$$c_t = c^* = K/T \tag{4}$$

Now assume that T is even and consider two small deviations from the optimal policy (cf. Figure 1)

$$C_t^1 = \begin{cases} c^* - \frac{1}{2}\epsilon & \text{for } t \text{ odd} \\ c^* + \frac{1}{2}\epsilon & \text{for } t \text{ even} \end{cases} \tag{5}$$

$$C_t^2 = c^* + \epsilon \left(t - \frac{T+1}{2} \right) \tag{6}$$

for small ϵ. Both policies satisfy the budget constraint (2). Up to a quadratic approxi-

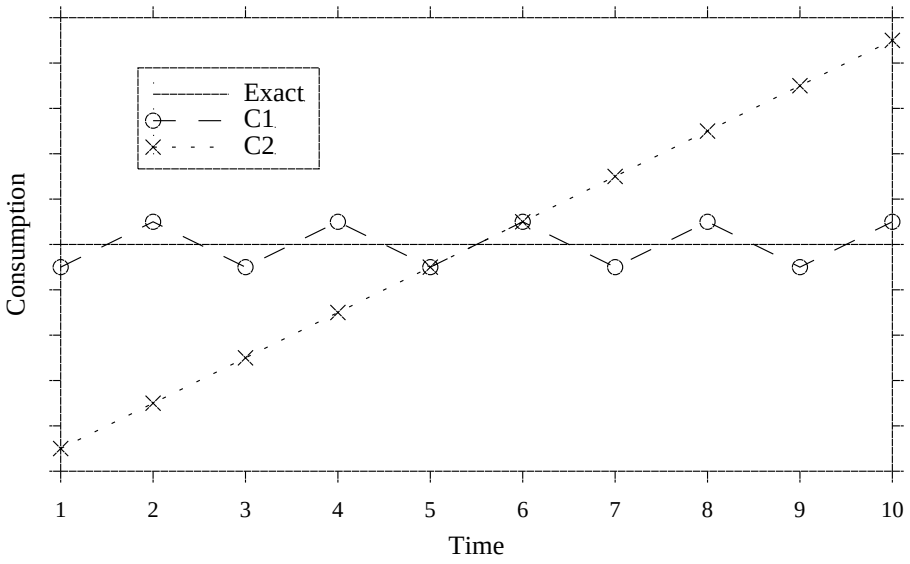

Figure 1: Example: finite horizon consumption problem

mation about c^*, the Euler residual for both policies and all t is given by $R_t = u''(c^*)\epsilon$, but the maximal approximation error of the policy function is 0.5ϵ for C_t^1 and $\frac{T-1}{2}\epsilon$ for C_t^2. The loss in value is $\frac{T}{8}|u''(c^*)|\epsilon^2$ for C_t^1 and $|u''(c^*)| \sum_{t=T/2+1}^{T} \epsilon^2 \left(t - \frac{T+1}{2}\right)^2$ for C_t^2.

The second policy has a much bigger error, because the mistake at a given point, as measured by the Euler residual, steadily adds up over time, while in the first policy the errors tend to cancel over time. This shows that there is no functional relationship between the maximum of the Euler residuals and the maximal value and policy error. A theory that relates maximal Euler residuals to approximation errors will in most cases substantially overestimate the error, since it has to account for the worst possible case (cf. the results in Santos, 2000, Section 4.3). If we want to obtain exact error estimates

based on Euler residuals, we therefore have to monitor the Euler residuals systematically over time. The present paper describes a method which does exactly this.

Note also that, with a given maximal Euler residual, the policy can deviate the more from the optimal policy, the longer the time horizon. In an infinite horizon problem, long time horizon should be interpreted as low discount factor, which explains the role of the discount factor in Santos (2000, Theorem 3.3).

3 The Model

We assume that the agent (or each agent in the economy) solves an optimization problems of the form

$$\max_{u_0,u_1,\ldots} \mathrm{E}_0 \sum_{t=0}^{\infty} \beta^t F(x_t, u_t, z_t) \tag{7a}$$

subject to

$$x_{t+1} = u_t + \xi(z_{t+1}), \qquad \mathrm{E}_t\, \xi(z_{t+1}) = 0 \tag{7b}$$

$$G(x_t, u_t, z_t) \geq 0 \tag{7c}$$

$$x_0 \text{ given} \tag{7d}$$

and the transversality condition

$$\mathrm{E}_0 \lim_{t\to\infty} \beta^t F_1(x_t, u_t, z_t) \cdot x_t = 0 \tag{7e}$$

Here, $x_t \in \Re^d$ is a vector of endogenous state variables, $u_t \in \Re^d$ is a vector of control variables, and z_t is a vector of exogenous random variables (in a general equilibrium model, z_t would also include all the aggregate variables that are exogenous to the individual decision maker). We assume that z_t is a Markovian process with transition function $Q(.,.)$ (defined as in Stokey and Lucas, 1989, p.212). The random vector can appear in the objective function as well as in the dynamic equation, with the requirement that $\mathrm{E}_t\, \xi(z_{t+1}) = 0$. In other words, the dynamic equation (7b) is written such that the control u_t is the expected value of next period's state vector, which simplifies notation. The function $F(x, u, z)$ is assumed to be measurable in z, and three times continuously differentiable in (x, u) for all z. We also assume that F is strictly concave in x and u.

The set of feasible controls $\Gamma(x, z) \equiv \{u | G(x, u, z) \geq 0\}$ is determined by the constraints (7c), where G is a function with values in $\Re^q$. We assume that $G(x, u, z)$ is concave in (x, u) for all z, which implies that the set of feasible combinations (x, u) is convex. Note that (7c) could be equivalently written in a way that G is only quasiconcave. For our purposes (cf. Section 4.4), however, it is necessary to write the constraints

in a form that G is really concave. The above assumptions guarantee that the Euler equation together with the transversality condition (7e) is sufficient for an optimum (Stokey and Lucas 1989, Theorem 4.15). For some results it is necessary to exclude increasing returns to scale by the assumption

$$\Gamma(\alpha x + (1-\alpha)x', z) \supseteq \alpha\Gamma(x,z) + (1-\alpha)\Gamma(x',z), \qquad \forall x, x', z, 0 \leq \alpha \leq 1 \tag{8}$$

This guarantees that the value function of problem (7) is concave in x (Stokey and Lucas 1989, Theorem 4.8).

We assume that the state variables (x, z) take values from a compact state space $\mathcal{Y}$. The continuity of the above functions and their derivatives, together with the compactness assumption implies that there is a constant $\overline{\nabla}$ such that $\|F_i\|, \|F_{ij}\|, \|G_i\|, \|G_{ij}\| \leq \frac{1}{2d}\overline{\nabla}$ for $i = 1, 2; j = 1, 2$. where $\|.\|$ is the supremum norm on vectors and matrices. This allows us to infer, for example, that

$$\|G(x+\delta, u+\epsilon, z) - G(x,u,z)\| \leq \overline{\nabla} \cdot \eta + O\left(\eta^2\right) \qquad \forall \delta, \epsilon : \|\delta\|, \|\epsilon\| < \eta \tag{9}$$

Compactness of the state space is adopted mainly to simplify proofs. What is critical is that the difference in utility between the true solution and the numerical solution along any simulated path grows at a rate lower than the discount rate. Compactness is a way to guarantee this: with $\mathcal{Y}$ compact, (7b) implies that the policy function $u(x,z)$ is bounded, and then F is bounded since it is continuous, and so the value function is bounded. While economic models are often defined on unbounded state spaces, numerical solutions have usually to be computed on a bounded set, often after a suitable transformation of the state variables.

Finally, one should note that (7) is not the most general stochastic optimization model, but the great majority of models studied in macroeconomics can be brought into this form (it is slightly more general than both model variants in Stokey and Lucas, 1989, p.239f.). This formulation will help to keep the notation as concise as possible, since the Euler equation then takes the simple form

$$R_{t+1} \equiv F_2\left(x_t, u_t, z_t\right) + \beta\, \mathrm{E}_t\, F_1\left(x_{t+1}, u_{t+1}, z_{t+1}\right) = 0 \tag{10}$$

As discussed above, the Euler residuals R_t of the numerical solution will form the basis for the estimation of the errors in the value and policy function.

4 Estimating the error in the value function

We assume that a numerical solution for the problem (7) is provided, and our task is to estimate the size of the approximation error of this solution. In this section we measure

accuracy as the decrease in the objective function that results from using the approximate rather than the exact solution. Section 5 will discuss the problem of measuring the error in the policy function.

4.1 Illustration: the deterministic case

This section discusses briefly and somewhat heuristically the deterministic case. This will provide some intuitive understanding of the accuracy estimates, even if we will later see that the analysis in the stochastic case differs substantially from the deterministic analysis.

Consider the deterministic version of the model with criterion function $F(x_t, u_t)$ and $x_{t+1} = u_t$. Assume we have simulated the path $\hat{x}_t$, for $t = 0, \ldots, T$, by the numerical solution, starting from $\hat{x}_0 = x_0$. Assume that the model is stable and that the simulated path converges at least approximately to the true stationary state. We can then choose T large enough so that what happens after time T is irrelevant, given a required degree of accuracy. The loss in value from using the numerical rather than the exact solution is therefore given, up to the required degree of accuracy, by the loss that results from choosing a suboptimal path from x_0 to $\hat{x}_T$. This loss is given by the solution to the problem

$$\max_{\delta_0, \delta_1, \ldots} \sum_{t=0}^{T} \beta^t F(\hat{x}_t + \delta_{t-1}, \hat{x}_{t+1} + \delta_t) - F(\hat{x}_t, \hat{x}_{t+1}) \tag{11a}$$

subject to (we abstract for the moment from inequality constraints)

$$\delta_T = 0 \tag{11b}$$

Note that

$$\delta_{-1} = 0 \tag{12}$$

since $x_0^* = \hat{x}_0 = x_0$ is given. The optimal δ_t's give us the policy error of the numerical solution, and the value function of this problem at $t = 0$ and x_0 gives us the value loss.

Of course, problem (11) is just a finite horizon version of problem (7) and equally hard to solve. However, if the numerical solution $\hat{x}_t$ is already close to the true solution, the above problem is approximated well by the locally quadratic approximation

$$\max_{\delta_0, \delta_1, \ldots} \sum_{t=0}^{T} \beta^t \left[\delta_{t-1}' F_1(\hat{x}_t, \hat{x}_{t+1}) + \delta_t' F_2(\hat{x}_t, \hat{x}_{t+1}) + \frac{1}{2} \left(\delta_{t-1}', \delta_t' \right) D^2 F \left(\hat{x}_t, \hat{x}_{t+1} \right) \begin{pmatrix} \delta_{t-1} \\ \delta_t \end{pmatrix} \right]$$

where F_1 and F_2 are the column vectors of partial derivatives w.r.t. the first and second argument, prime denotes transposition and $D^2 F$ is the Hessian of F. This problem can be solved by backward induction, with little numerical effort. The procedure gives an estimate of the policy error δ as well as the value loss from the approximation at x_0.

(By iterating the quadratic approximation, we would probably converge to the exact solution, but this is not our concern here.)

From this we learn that, if we already have a reasonably good numerical solution, evaluating its accuracy is a much simpler task then finding the solution in the first place. The reason is that we can use the information that is contained in the solution, which gives us the realized paths about which we can locally approximate. Obviously, the precision with which we can measure the accuracy of the solution depends itself on the accuracy of the solution: the smaller the size of the δ_t's, the better the locally quadratic approximation. This will be made more precise in the next section. Since economic problems usually work with very smooth functions, we can expect to obtain good estimates of the precision even if the numerical solution is rather weak. The numerical examples of Section 6 will confirm this.

4.2 Stochastic model: an upper bound

From now on we deal with the stochastic model of Section 3. In this and the next subsection, we abstract from inequality constraints and assume that the exact as well as the proposed numerical solution are inner solutions to this problem. Inequality constraints will be dealt with in Subsection 4.4.

Again, we follow the idea of estimating the value and policy error along realized paths, but the argument is now more complex, for the following reason. Solving the deterministic program (11) backward in time, we can compute δ_t if we know next period's value function in a neighborhood of $\hat{x}_{t+1}$. This problem can be well approximated then by a locally quadratic approximation about the $(\hat{x}_t, \hat{x}_{t+1})$, since the approximation needs to hold only over a range of values of the order of the δ_t's. In contrast, to find the δ_t in a stochastic problem, we need to know the value function at all possible realizations of x_{t+1} that start from $\hat{x}_t$. Since the shocks can be large, we need the value function in a big part of the state space, and a quadratic approximation cannot be accurate in general. In particular, one cannot carry over the analysis of the deterministic case to the stochastic model by solving the stochastic version of the quadratic approximation about the realized path. It is therefore not possible to compute the value loss with the same precision as in the deterministic case. However, the analysis of this section will provide a tight *upper bound* to the value loss which is based on purely *local* quadratic approximations.

We first introduce some notation. A realized path of the optimal solution will always be denoted by (x_t^*, u_t^*), a path of the numerical solution by $(\hat{x}_t, \hat{u}_t)$. These paths depends on the (given) initial value x_0 and a realization of the shocks $z_0, z_1, \ldots$. Let us denote the history of shocks $z_0, z_1, \ldots, z_t$ by z^t. Because of the Markov structure of the problem,

the optimal policy u_t^* is a function of the vector of state variables, $u_t^* = U^*(x_t^*, z_t)$. We also assume that the policy $\hat{u}_t = \hat{U}(\hat{x}_t, z_t)$ provided by the numerical solution is a function of the state vector. This also implies that, with a fixed initial state x_0, the realizations x_t^*, u_t^*, $\hat{x}_t$ and $\hat{u}_t$ are a function of the history of shocks z^t. All expectation operators should be understood as referring to the distribution of the shocks $z_0, z_1, \ldots$. Formally, the operator E_t means expectation w.r.t. the σ-algebra generated by z^t. We will say that a random variable is t-measurable to express that it is a function of z^t. We define the policy error $\delta_t(z^t) \equiv u_t^*(z^t) - \hat{u}_t(z^t) = U^*(x_t^*, z_t) - \hat{U}(\hat{x}_t, z_t)$. The notation stresses the fact that δ_t has to be understood as a function of the history of shocks z^t, not of the current state vector. From (7b) it follows that the optimal state x_t^* satisfies $x_t^* = \hat{x}_t + \delta_{t-1}$.

Using this notation, we can now define the approximation error Δ of the value function as

$$\Delta \equiv \mathrm{E}_0 \sum_{t=0}^{\infty} \beta^t \Big[F(x_t^*, u_t^*, z_t) - F(\hat{x}_t, \hat{u}_t, z_t) \Big]$$

We truncate the error term after T periods and define

$$\begin{aligned} \Delta^{trunc} &\equiv \mathrm{E}_0 \sum_{t=0}^{T} \beta^t \Big[F(x_t^*, u_t^*, z_t) - F(\hat{x}_t, \hat{u}_t, z_t) \Big] \\ &= \mathrm{E}_0 \sum_{t=0}^{T} \beta^t \left[F(\hat{x}_t + \delta_{t-1}, \hat{u}_t + \delta_t, z_t) - F(\hat{x}_t, \hat{u}_t, z_t) \right] \end{aligned} \tag{13}$$

where $\delta_{-1} = 0$ by construction. The difference between Δ and Δ^{trunc} is given by $\|\Delta - \Delta^{trunc}\| = \mathrm{E}_0 \beta^{T+1} \| V(x_{T+1}^*) - W(\hat{x}_{T+1}) \|$, where $V(x)$ is the (exact) value function at point x, and $W(x)$ is the value reached under the policy of the numerical solution. The boundedness of F shows that $\|\Delta - \Delta^{trunc}\|$ goes to zero for $T \to \infty$, so we can concentrate on measuring Δ^{trunc} rather than Δ if we choose T large enough.

As noted above, the precision of our accuracy estimates will itself depend on the accuracy of the numerical solution, and on the time horizon for which we simulate. To make this precise we define

$$\eta \equiv \max \left\{ \sup_{t, z^t} \|\beta^t \delta_t(z^t)\|, \beta^{T/3} \right\} \tag{14}$$

Note that η is finite since the policy functions are bounded. We will show that the true error Δ is of order $O(\eta^2)$. In estimating the upper bound, we will rely on some Taylor approximations, whose effect on the estimate will be shown to be of order $O(\eta^3)$. Let us first give a precise definition of the order terms. Assume that some variable χ is a function of the numerical solution. To say that χ is of order $O(\eta^m)$ then means

the following. For any sequence of numerical solutions, indexed by n, such that the corresponding sequence of supremum errors η_n goes to 0 for $n \to \infty$, $|\chi/\eta_n^m| \leq M$ for some constant M. To say that χ is of order $o(\eta^m)$ means that, for any such sequence, $\lim_{n\to\infty} |\chi/\eta_n^m| = 0$. Note that with the above definition of η we can write $\Delta = \Delta^{trunc} + O(\eta^3)$.

To estimate the upper bound, we proceed in three steps. First, we derive a simple expression for Δ^{trunc}, based on a quadratic approximation. Second we show a way to compute an upper bound to this quadratic approximation based on an estimate of the policy error. Third, we show how to obtain the policy error estimate.

Step 1: Quadratic approximation to Δ^{trunc}

We start with a second order approximation of (13) about $(\hat{x}_t, \hat{u}_t)$. Since the third partial derivatives of F are bounded, Taylor's theorem shows that

$$\begin{aligned}\Delta^{appr} \equiv \mathrm{E}_0 \sum_{t=0}^{T} \beta^t \Big[& \delta_{t-1}' F_{1,t}(\hat{x}_t, \hat{u}_t, z_t) + \delta_t' F_{2,t}(\hat{x}_t, \hat{u}_t, z_t) \\ & + \frac{1}{2}\Big(\delta_{t-1}' F_{11,t}(\hat{x}_t, \hat{u}_t, z_t)\,\delta_{t-1} + 2\delta_{t-1}' F_{12,t}(\hat{x}_t, \hat{u}_t, z_t)\,\delta_t + \delta_t' F_{22,t}(\hat{x}_t, \hat{u}_t, z_t)\,\delta_t\Big)\Big] \\ & = \Delta^{trunc} + O(\eta^3) \end{aligned} \quad (15)$$

Using the definition

$$Quad(F, a, b; t) \equiv a_{t-1}' F_{11,t} b_{t-1} + a_{t-1}' F_{12,t} b_t + a_t' F_{21,t} b_{t-1} + a_t' F_{22,t} b_t \quad (16)$$

for any function F and series of vectors a and b, the following lemma derives a simpler expression for Δ^{appr}. (When the arguments of the derivatives $F_{1,t}$ etc. are omitted, it should always be understood that they are taken at the simulated points $(\hat{x}_t, \hat{u}_t, z_t)$.)

Lemma 1.

$$\Delta^{appr} = -\frac{1}{2}\mathrm{E}_0 \sum_{t=0}^{T} \beta^t Quad(F, \delta, \delta; t) + O(\eta^3) \quad (17)$$

Proof. Along any path z^{T+1}, the approximation errors δ_t satisfy the Euler equation

$$F_2(\hat{x}_t + \delta_{t-1}, \hat{u}_t + \delta_t, z_t) + \beta\,\mathrm{E}_t\, F_1(\hat{x}_{t+1} + \delta_t, \hat{u}_{t+1} + \delta_{t+1}, z_{t+1}) = 0 \quad (18)$$

Using again Taylor's theorem, (18) can be approximated as

$$F_{2,t} + F_{21,t}\delta_{t-1} + F_{22,t}\delta_t + \beta\,\mathrm{E}_t\,(F_{1,t+1} + F_{11,t+1}\delta_t + F_{12,t+1}\delta_{t+1}) = \quad (19)$$

$$R_{t+1} + F_{21,t}\delta_{t-1} + (F_{22,t} + \beta\,\mathrm{E}_t\,F_{11,t+1})\,\delta_t + \beta\,\mathrm{E}_t\,[F_{12,t+1}\delta_{t+1}] = O(\eta^2) \quad (20)$$

If we premultiply (19) by δ_t', use the fact that δ_t is a function of z^t so that $\delta_t' \mathrm{E}_t(\chi) = \mathrm{E}_t[\delta_t'\chi]$ (Billingsley 1986, Theorem 34.3) for any vector of random variables χ, take expected values and use the law of iterated expectations, we get

$$\mathrm{E}_0\Big[\delta_t' F_{2,t} + \delta_t' F_{21,t}\delta_{t-1} + \delta_t' F_{22,t}\delta_t + \beta\left(\delta_t' F_{1,t+1} + \delta_t' F_{11,t+1}\delta_t + \delta_t' F_{12,t+1}\delta_{t+1}\right)\Big] = O\left(\eta^3\right) \quad (21)$$

Multiplying (21) by β^t, summing up for $t = 0, \ldots, T-1$ and using $\delta_{-1} = 0$, we get

$$\mathrm{E}_0 \sum_{t=0}^{T} \beta^t \Big[\delta_{t-1}' F_{1,t} + \delta_t' F_{2,t} + \delta_{t-1}' F_{11,t}\delta_{t-1} + 2\delta_{t-1}' F_{12,t}\delta_t + \delta_t' F_{22,t}\delta_t\Big] = \mathrm{E}_0\, \beta^T \Big[F_{2,T}\delta_T + \Big(\delta_T' F_{21,T}\delta_{T-1} + \delta_T' F_{22,T}\delta_T\Big)\Big] + O\left(\eta^3\right) = O\left(\eta^3\right) \quad (22)$$

where we have used $\beta^T < \eta^3$ (cf. (14)) as well as the following reordering of terms:

$$\sum_{t=0}^{T} \beta^t \left(\delta_{t-1}' F_{11,t}\delta_{t-1} + 2\delta_{t-1}' F_{12,t}\delta_t + \delta_t' F_{22,t}\delta_t\right) = \sum_{t=0}^{T-1} \beta^t \left[\delta_t' F_{21,t}\delta_{t-1} + \delta_t' F_{22,t}\delta_t + \beta\left(\delta_t' F_{11,t+1}\delta_t + \delta_t' F_{12,t+1}\delta_{t+1}\right)\right] + \delta_{-1}' F_{11,0}\delta_{-1} + \delta_{-1}' F_{12,0}\delta_0 + \beta^T\left(\delta_T' F_{21,T}\delta_{T-1} + \delta_T' F_{22,T}\delta_T\right) \quad (23)$$

Using (22) in (15) we obtain (17). □

Our task therefore reduces to estimating the quadratic term in (17), which is non-negative, since the concavity assumption on F means that $Quad\,(F, \delta, \delta; t)$ is a negative definite quadratic form.

Step 2: Computing an upper bound for Δ^{appr} using an estimate of δ_t

As discussed at the beginning of this section, we cannot "compute" the δ_t's in the stochastic model, not even for a given realization of z^T. This becomes clear from Equ. (20): the Euler residual R_{t+1} can be measured, but the expectation $\beta\,\mathrm{E}_t[F_{12,t+1}\delta_{t+1}]$ cannot, because it requires the values of δ_{t+1} off the realized path. Computing δ_0 to δ_T recursively would require the values of δ at all realizations starting from x_0.

However, we will see below that we can use an equation similar to (20) to obtain an estimate of δ_t, which we denote by $\bar{\delta}_t$. If we compute (17) with $\bar{\delta}_t$ rather than δ_t, we can define

$$\Delta^{ub} \equiv -\frac{1}{2}\,\mathrm{E}_0 \sum_{t=0}^{T} \beta^t Quad\,(F, \bar{\delta}, \bar{\delta}; t) \quad (24)$$

Denoting the error in estimating δ_t by $\tilde{\delta}_t \equiv \delta_t - \bar{\delta}_t$, we get from (17)

$$\Delta^{ub} = -\frac{1}{2}\,\mathrm{E}_0 \sum_{t=0}^{T} \beta^t Quad\left(F, \delta - \tilde{\delta}, \delta - \tilde{\delta}; t\right) \tag{25}$$

$$= \Delta^{appr} - 2\Delta^{cross} + \Delta^{esterr} + O\left(\eta^3\right) \tag{26}$$

where

$$\Delta^{cross} \equiv -\frac{1}{2}\,\mathrm{E}_0 \sum_{t=0}^{T} \beta^t Quad\left(F, \delta, \tilde{\delta}; t\right) \tag{27}$$

$$\Delta^{esterr} \equiv -\frac{1}{2}\,\mathrm{E}_0 \sum_{t=0}^{T} \beta^t Quad\left(F, \tilde{\delta}, \tilde{\delta}; t\right) \tag{28}$$

Since $F(x, u, z)$ is concave in (x, u), we have $\Delta^{esterr} > 0$, and we reach the conclusion that $\Delta^{ub} \geq \Delta^{appr} + O\left(\eta^3\right)$ if we can show that

$$\Delta^{cross} = O\left(\eta^3\right) \tag{29}$$

In other words, under this condition, Δ^{ub} is an upper bound for Δ^{appr}, up to a cubic error term (note from (17) that Δ^{appr} is of order $O\left(\eta^2\right)$). We can compute Δ^{ub} from (24) by Monte-Carlo techniques, to any desired degree of accuracy.

Step 3: Estimating δ_t

We still have to find an estimate $\bar{\delta}_t$ of δ_t such that condition (29) is satisfied. To do this we start from (19). While the term $\mathrm{E}_t\, F_{1,t+1}$, the expectational part of the Euler equation, can in principle be measured with arbitrary precision, we allow for generality that it be replaced by an estimate $F^e_{1,t+1}$ that has the property $\mathrm{E}_t\, F^e_{1,t+1} = \mathrm{E}_t\, F_{1,t+1}$. Any estimate of the expectation obtained by Monte-Carlo integration has this property. As discussed above, the term $\mathrm{E}_t\left[F_{11,t+1}\delta_t + F_{12,t+1}\delta_{t+1}\right]$ is not available, because we have no estimate of δ_{t+1} off the realized path. In the following we simply ignore the expectation operator E_t and use $\left[F_{11,t+1}\bar{\delta}_t + F_{12,t+1}\bar{\delta}_{t+1}\right]$ (note that the realization can be interpreted as an unbiased estimate of its expected value). Then we obtain the second order difference equation in $\bar{\delta}_t$

$$F_{2,t} + F_{21,t}\bar{\delta}_{t-1} + F_{22,t}\bar{\delta}_t + \beta\left(F^e_{1,t+1} + F_{11,t+1}\bar{\delta}_t + F_{12,t+1}\bar{\delta}_{t+1}\right) = 0 \tag{30}$$

Imposing the boundary conditions

$$\bar{\delta}_{-1} = \bar{\delta}_T = 0, \tag{31}$$

Equ. (30) determines $\bar{\delta}_t$ for any given realization of z^t. An efficient way to compute the $\bar{\delta}_t$ recursively is provided in Section 4.3.

Lemma 2. *If $\bar{\delta}_t$ is defined by Equ. (30) and the boundary conditions (31), then Equ. (29) holds.*

Proof. Subtracting (30) from (19) and conditioning on t we get

$$\mathrm{E}_t\left[F_{21,t}\tilde{\delta}_{t-1} + F_{22,t}\tilde{\delta}_t + \beta\left(F_{11,t+1}\tilde{\delta}_t + F_{12,t+1}\tilde{\delta}_{t+1}\right)\right] = O\left(\eta^2\right) \tag{32}$$

Premultiplying (32) by δ_t', using again $\delta_t' \mathrm{E}_t(\chi) = \mathrm{E}_t\left[\delta_t'\chi\right]$, and taking expectations E_0, we get

$$\mathrm{E}_0\left[\delta_t'F_{21,t}\tilde{\delta}_{t-1} + \delta_t'F_{22,t}\tilde{\delta}_t + \beta\left(\delta_t'F_{11,t+1}\tilde{\delta}_t + \delta_t'F_{12,t+1}\tilde{\delta}_{t+1}\right)\right] = O\left(\eta^3\right) \tag{33}$$

If we multiply (33) by β^t, sum up for $t = 0, \ldots, T-1$, and use the reordering (23), we obtain

$$\mathrm{E}_0 \sum_{t=0}^{T} \beta^t Quad\left(F, \delta, \tilde{\delta}; t\right) = \beta^T\left(\delta_T'F_{21,T}\tilde{\delta}_{T-1} + \delta_T'F_{22,T}\tilde{\delta}_T\right) + O\left(\eta^3\right) = O\left(\eta^3\right) \tag{34}$$

□

Note the following two properties of the estimate $\bar{\delta}_t$:

1. If the estimated Euler residual is zero for all t, Equ. (30) is a homogeneous difference equation, and the boundary conditions (31) imply that the estimated error $\bar{\delta}_t$ is zero for all t.

2. $\bar{\delta}_t$ depends on the whole realization $z_0, \ldots, z_T$, so $\bar{\delta}_t$ is not t-measurable.

The calculations of this section are summarized by the following theorem.

Theorem 1. *If $\bar{\delta}_t$ is defined by Equ. (30) and the boundary conditions (31), then*

$$\Delta^{ub} \geq \Delta + O\left(\eta^3\right) \tag{35}$$

Theorem 1 is the first main result of this paper. It shows that even in the stochastic case, the estimated errors along realized paths can be used to estimate an upper bound of the approximation error of the value function. The numerical results of Section 6.2 will show that this upper bound is rather tight.

It may be useful to list all sources of error that come into play in estimating Δ by Δ^{ub}:

1. The error from the quadratic approximation of F and the linear approximation of the Euler equation about $(\hat{x}_t, \hat{u}_t)$, which together give an error which is $O\left(\eta^3\right)$.

2. The truncation error from the finite horizon of the simulation.

3. The sampling error in the simulation.

4. Imprecise measurement of the Euler residual.

5. The quadratic term Δ^{esterr}.

The errors 2, 3 and 4 can be made arbitrarily small by increasing the computational effort, using longer time horizons, more simulations etc. The error 5 cannot be decreased by doing more simulations etc., and is quadratic in the maximal deviation, just as Δ^{ub} itself. However, we know that it is positive, so we can say that Δ^{ub} is an upper bound. The error 1 depends on the nonlinearity of the model. It is a cubic function of the maximal deviation of the simulation, so it disappears fast when the numerical solution goes to the exact solution. For a given solution, however, it cannot be decreased, and if the numerical solution is very imprecise, the cubic term may dominate other terms and make that Δ^{ub} is not really an upper bound.

4.3 Recursive computation of the error estimate

This section provides recursive formulas for the computation of the quadratic term in (24) for a given realization, that means, for a given history of shocks z^T. The formulas are similar to the recursions of the linear-quadratic control problem. If we define

$$R^e_{t+1} = F_{2,t} + \beta F^e_{1,t+1} \tag{36}$$

we have from (30)

$$\bar{\delta}_t = -\left(F_{22,t} + \beta F_{11,t+1}\right)^{-1} \left(F_{21,t}\bar{\delta}_{t-1} + R^e_{t+1} + \beta F_{12,t+1}\bar{\delta}_{t+1}\right) \tag{37}$$

Since $\bar{\delta}_T = 0$ from (31), we get

$$\bar{\delta}_{T-1} = a_{T-1} + A_{T-1}\bar{\delta}_{T-2} \tag{38a}$$

where

$$a_{T-1} \equiv -\left(F_{22,T-1} + \beta F_{11,T}\right)^{-1} R^e_T \tag{38b}$$

$$A_{T-1} \equiv -\left(F_{22,T-1} + \beta F_{11,T}\right)^{-1} F_{21,T-1} \tag{38c}$$

This can be used as the starting point of the following recursion. If $\bar{\delta}_{t+1} = a_{t+1} + A_{t+1}\bar{\delta}_t$ is given with some vector a_{t+1} and some matrix A_{t+1}, it follows from (37) that

$$\bar{\delta}_t = a_t + A_t\bar{\delta}_{t-1} \tag{39a}$$

with

$$a_t = -\left(F_{22,t} + \beta F_{11,t+1} + \beta F_{12,t+1}A_{t+1}\right)^{-1} \left(R^e_{t+1} + \beta F_{12,t+1}a_{t+1}\right) \tag{39b}$$

$$A_t = -\left(F_{22,t} + \beta F_{11,t+1} + \beta F_{12,t+1}A_{t+1}\right)^{-1} F_{21,t} \tag{39c}$$

Iterating (39a) forward, we obtain $\bar{\delta}_{t+s}$ as a function of $\bar{\delta}_t$ for $s > 0$. This allows us to define

$$\Delta_t^{ub}(\bar{\delta}_{t-1}; z^T) \equiv -\frac{1}{2}\sum_{s=t}^{T} \beta^{s-t} \left(\bar{\delta}_{s-1}\left(\bar{\delta}_{t-1}\right)', \bar{\delta}_s\left(\bar{\delta}_{t-1}\right)'\right) D^2 F_s \begin{pmatrix} \bar{\delta}_{s-1}\left(\bar{\delta}_{t-1}\right) \\ \bar{\delta}_s\left(\bar{\delta}_{t-1}\right) \end{pmatrix} \tag{40}$$

and compute it recursively, starting with

$$\Delta_T^{ub}(\bar{\delta}_{T-1}; z^T) = -\frac{1}{2}\bar{\delta}'_{T-1} F_{11,T} \bar{\delta}_{T-1} \tag{41}$$

If we have

$$-2\Delta_{t+1}^{ub}(\bar{\delta}_t; z^T) = \bar{v}_{t+1} + v'_{t+1}\bar{\delta}_t + \bar{\delta}'_t V_{t+1}\bar{\delta}_t \tag{42}$$

for some given constant $\bar{v}_{t+1}$, vector v_{t+1} and matrix V_{t+1}, it follows from (39) that

$$-2\Delta_t^{ub}(\bar{\delta}_{t-1}; z^T) = \bar{v}_t + v'_t\bar{\delta}_{t-1} + \bar{\delta}'_{t-1} V_t \bar{\delta}_{t-1} \tag{43a}$$

with

$$\bar{v}_t = a'_t F_{22,t} a_t + \beta\left[\bar{v}_{t+1} + v'_{t+1}a_t + a'_t V_{t+1} a_t\right] \tag{43b}$$

$$v_t = 2F_{12,t}a_t + 2A'_t F_{22,t} a_t + \beta\left[A'_t v_{t+1} + 2A'_t V_{t+1} a_t\right] \tag{43c}$$

$$V_t = F_{11,t} + 2F_{12,t}A_t + A'_t F_{22,t} A_t + \beta A'_t V_{t+1} A_t \tag{43d}$$

Since $\bar{\delta}_{-1} = 0$, we finally arrive at

$$-\frac{1}{2}\sum_{t=0}^{T} \beta^t Quad\left(F, \bar{\delta}, \bar{\delta}; t\right) = \Delta_0^{ub}(0; z^T) = -\frac{1}{2}\bar{v}_0 \tag{44}$$

The estimate Δ^{ub} in (25) is the average of $\Delta_0^{ub}(0; z^T)$ over many simulated paths z^T. To compute the error at a specific point x_0 of the state space, the simulations will all start from x_0. The compute the (weighted) average of the error over the state space, we start the simulations from different points, drawn from a suitable distribution over the state space.

4.4 Inequality constraints

The following paragraphs show how the analysis of the last sections has to be modified to account for occasionally binding inequality constraints of the form (7c).

Equ. (7c) defines $q > 0$ constraints. Constraint i, $i = 1, \ldots, q$ will be written $G^i(x_t, u_t, z_t) \geq 0$ and will be related to a Lagrange multiplier λ_t^i, which is a function of the state variables (x_t, z_t), satisfying the usual complementary slackness conditions $\lambda_t^i \geq 0$ and $\lambda_t^i G_t^i = 0$. It is important to bear in mind that the G^i are convex, not just quasi-convex functions. Quasi-convex restrictions can of course always be brought into an equivalent convex form.

In the new framework, the Euler equation is

$$0 = F_2\left(x_t^*, u_t^*, z_t\right) + \sum_{i=1}^{q} G_2^i\left(x_t^*, u_t^*, z_t\right) \lambda_t^i(x_t^*, z_t)$$
$$+ \beta\, \mathrm{E}_t\left[F_1\left(x_{t+1}^*, u_{t+1}^*, z_{t+1}\right) + \sum_{i=1}^{q} G_1^i\left(x_{t+1}^*, u_{t+1}^*, z_{t+1}\right) \lambda_{t+1}^i\left(x_{t+1}^*, z_{t+1}\right)\right] \quad (45)$$

Since the Euler equation depends on the Lagrange multipliers, it is necessary to have numerical estimates $\hat{\lambda}^i\left(x_t, z_t\right)$ available. This should not cause any problem, since most methods (for example projection methods) compute the Lagrange multiplier function as an element of the solution. In a dynamic programming framework, the multipliers are not needed, but are easy to obtain once a solution has been computed. Using the estimated Lagrange multipliers, the measured Euler residual is

$$R_{t+1} \equiv F_2\left(\hat{x}_t, \hat{u}_t, z_t\right) + \sum_{i} G_2^i\left(\hat{x}_t, \hat{u}_t, z_t\right) \hat{\lambda}_t^i(\hat{x}_t, z_t)$$
$$+ \beta\, \mathrm{E}_t\left[F_1\left(\hat{x}_{t+1}, \hat{u}_{t+1}, z_{t+1}\right) + \sum_{i} G_1^i\left(\hat{x}_{t+1}, \hat{u}_{t+1}, z_{t+1}\right) \hat{\lambda}_{t+1}^i\left(\hat{x}_{t+1}, z_{t+1}\right)\right] \quad (46)$$

We impose the following regularity conditions on the numerical policy function and Lagrange multipliers:

Assumption 1. *The approximate policy* $\hat{U}(.)$ *and the estimated Lagrange multipliers* $\hat{\lambda}^i(.)$

i) are both functions of the state variables (x, z),

ii) satisfy for all x, z *and* i *the complementary slackness conditions*

$$\hat{\lambda}^i(x, z) \geq 0, \qquad G^i(x, \hat{U}\left(x, z\right), z) \geq 0, \qquad \hat{\lambda}^i(x, z) G^i(x, \hat{U}\left(x, z\right), z) = 0$$

The error bound Δ^{ub} will now depend not only on the policy error δ_t, but also on the properties of the estimated Lagrange multipliers. We therefore update our definition of the supremum error to

$$\eta_c \equiv \max\left\{\sup_{t, z^t, i} \|\beta^t \delta_t(z^t)\|, \sup_{(x,z)\in\mathcal{Y}} \|\hat{\lambda}^i(x, z) - \lambda^i(x, z)\|, \beta^{T/3}\right\} \quad (47)$$

Using the abbreviations

$$\hat{H}_{jk,t} \equiv F_{jk}\left(\hat{x}_t, \hat{u}_t, z_t\right) + \sum_{i=1}^{q} \hat{\lambda}^i(\hat{x}_t, z_t) G_{jk}^i\left(\hat{x}_t, \hat{u}_t, z_t\right), \qquad j, k = 1, 2$$
$$Quad\left(\hat{H}, a, b; t\right) \equiv a_{t-1}' \hat{H}_{11,t} b_{t-1} + a_{t-1}' \hat{H}_{12,t} b_t + a_t' \hat{H}_{21,t} b_{t-1} + a_t' \hat{H}_{22,t} b_t$$

the estimate $\bar{\delta}_t$ of the policy error will now be calculated by

$$R^e_{t+1} + \hat{H}_{21,t}\bar{\delta}_{t-1} + \hat{H}_{22,t}\bar{\delta}_t + \beta\left[\hat{H}_{11,t+1}\bar{\delta}_t + \hat{H}_{12,t+1}\bar{\delta}_{t+1}\right] = 0 \tag{48}$$

which is analogous to (30), and the value loss will be estimated by

$$\Delta^{ub\_constr} \equiv -\frac{1}{2}\,\mathrm{E}_0 \sum_{t=0}^{T} \beta^t Quad\left(\hat{H}, \bar{\delta}, \bar{\delta}; t\right) \tag{49}$$

Again, R^e_{t+1} is any estimate that satisfies $\mathrm{E}_t\, R^e_{t+1} = R_{t+1}$. Our aim is to show that the statistic Δ^{ub_constr} provides an upper bound, up to an $O\left({\eta_c}^3\right)$ error term, to the quadratic approximation Δ^{appr} of the value loss, which is still given by (15). This is accomplished by the following lemma (the proofs of all lemmas are given in Appendix A).

Lemma 3.

$$\Delta^{ub\_constr} = \Delta^{appr} + \Delta^{esterr\_c} + \mathrm{E}_0 \sum_{t=0}^{T} \beta^t \sum_{i=1}^{q} G^i\left(x^*_t, u^*_t, z_t\right) \hat{\lambda}^i_t(\hat{x}_t, z_t) + O\left({\eta_c}^3\right) \tag{50}$$

where

$$\Delta^{esterr\_c} \equiv -\frac{1}{2}\,\mathrm{E}_0 \sum_{t=0}^{T} \beta^t Quad\left(\hat{H}, \tilde{\delta}, \tilde{\delta}; t\right)$$

Since $\hat{H}$ is negative definite by our assumptions on F and G, and $G^i\left(x^*_t, u^*_t, z_t\right) \geq 0$ and $\hat{\lambda}^i_t \geq 0$, we obtain the following theorem, which is analogous to Theorem 1.

Theorem 2. *If $\bar{\delta}_t$ is defined by Equ. (48) and the boundary conditions (31), then*

$$\Delta^{ub\_constr} \geq \Delta + O\left({\eta_c}^3\right)$$

The theorem provides the desired upper bound. However, to have *tight* bounds, it is of interest to know whether the term $\mathrm{E}_0 \sum_{t=0}^{T} \beta^t \sum_{i=1}^{q} G^i\left(x^*_t, u^*_t, z_t\right) \hat{\lambda}^i_t$ in (50) is big compared to Δ^{appr} or to Δ^{esterr_c}. Intuitively, we would expect it to be small, since the term is nonzero only in cases where the constraint is binding in the numerical solution, but not in the exact solution (recall that x^*_t, $\hat{x}_t$, $\hat{\lambda}_t$ etc. are all functions of the history of shocks z^t), which should not happen very often if the numerical solution is good. The following two lemmas make this precise. They show that the additional error term that appears in the case of inequality constraints vanishes quickly when the numerical solution becomes more accurate.

Lemma 4. *Assume that the approximate policy function $\hat{U}(.)$ and the estimated Lagrange multipliers $\hat{\lambda}^i(.)$ are locally Lipschitz, in the sense that there exist constants Λ and $\epsilon_0 > 0$ such that for all x, z, i, and $\epsilon < \epsilon_0$*

$$\left\|(x,z) - (x^1,z^1)\right\| < \epsilon \;\Rightarrow\; \left\|\hat{\lambda}^i(x,z) - \hat{\lambda}^i(x^1,z^1)\right\| < \Lambda\epsilon \;\; and \;\; \left\|\hat{U}(x,z) - \hat{U}\left(x^1,z^1\right)\right\| < \Lambda\epsilon \tag{51}$$

Then, under Assumption 1,

$$\mathrm{E}_0 \sum_{t=0}^{T} \beta^t \sum_{i=1}^{q} G^i\left(x_t^*, u_t^*, z_t\right) \hat{\lambda}_t^i(\hat{x}_t, z_t) = o\left({\eta_c}^2\right) \tag{52}$$

Lemma 5. *In addition to (51), assume that there exist constants $\bar{P}_\lambda$ and ϵ_0 such that for all i, t and $\epsilon < \epsilon_0$*

$$Pr\left\{\hat{\lambda}^i(\hat{x}_t, z_t) \in (0, \epsilon]\right\} \leq \bar{P}_\lambda \epsilon \tag{53}$$

where $(a, b]$ denotes the half-open interval excluding a and including b. Then

$$\mathrm{E}_0 \sum_{t=0}^{T} \beta^t \sum_{i=1}^{q} G^i\left(x_t^*, u_t^*, z_t\right) \hat{\lambda}_t^i(\hat{x}_t, z_t) = O\left({\eta_c}^3\right) \tag{54}$$

The assumption in Lemma 5 basically says that, close to the region where the constraint i is exactly binding, the probability density of $\hat{\lambda}^i$ is bounded. For a given numerical solution, Equ. (53) is easy to check numerically, but it is more difficult to establish general conditions that lead to (53). Appendix B therefore provides a more detailed discussion of this assumption and shows that it actually holds under quite natural conditions.

The conclusion of the above discussion is that the analysis of Section 4.2 and the recursive computation of Section 4.3 go through in the case of occasionally binding constraints. We only have to replace the terms F_{ij} by $\hat{H}_{ij}$.[1]

5 Estimating the error in the policy function

The error in the policy function is more difficult to measure than the error in the value function. In the procedure of the last section, we actually computed an estimate $\bar{\delta}_t$ of the policy error, but there is no theorem showing that this estimate were unbiased or an upper bound of the true policy error. In estimating the policy error we make a mistake, which we denoted by $\tilde{\delta}_t$, and everything we know about it is Equ. (32). The main obstacle to further analysis is that we cannot say anything about the correlation between $\tilde{\delta}_{t+1}$ and $F_{12,t+1}$ etc. If we knew that they were uncorrelated, we could take unconditional expectations in (32), replace $\mathrm{E}_0[F_{11,t}\tilde{\delta}_{t+1}]$ by $\mathrm{E}_0 F_{11,t}\,\mathrm{E}_0\,\tilde{\delta}_{t+1}$ etc., and obtain a homogeneous second order difference equation in $\mathrm{E}_0\,\tilde{\delta}_t$. Since $\mathrm{E}_0\,\tilde{\delta}_{-1} = \mathrm{E}_0\,\tilde{\delta}_T = 0$, we could then infer $\mathrm{E}_0\,\tilde{\delta}_t = 0$ for all t. In the limiting case where the Hessian is constant, this reasoning goes through. It also suggests that for smooth problems, where the Hessian moves slowly over time, the estimate is almost unbiased. In the general

[1] Note that in the case of linear inequality constraints, the G_{ij} vanish and $F_{ij} = \hat{H}_{ij}$. The Lagrange multipliers are nevertheless needed to compute the Euler residuals, cf. Equ. (46).

case, however, since the policy error is an arbitrary function of the state variables, it is not clear whether this is approximately fulfilled or not.

In the general case, one can then follow three strategies:

1. If we only want to know the approximate size of the policy errors and do not need rigorous error bounds, we can estimate it by the mean absolute value of the estimated policy error. The numerical experiments with variants of the stochastic growth model (cf. Section 6.1) suggest that it is a good estimate for well-behaved problems.

2. If the value function is strongly concave, the upper bound on the value loss can be used to obtain an upper bound on the expected squared policy error. This will be treated in Section 5.1.

3. If one is willing to make a higher computational effort than that of Section 4.2 (but still less than the one needed to obtain the solution of the optimization problem), it is possible to compute the policy error at a certain point up to $O\left({\eta_c}^2\right)$. This will be shown in Section 5.2.

5.1 Bounds on the average policy error

Having computed a bound on the value function loss, one can obtain bounds on the policy errors if the value function is strongly concave in a sense to be specified below.

To obtain tight error bounds, the argument of this section exploits the idea that, in a stationary solution, on average only the fraction $(1-\beta)$ of the value error is attributable to the current policy error, the rest being due to future policy errors. The following paragraphs will make this precise.

Denote by $V(x,z)$ the true value function, and by $W(x,z)$ the value that is obtained by following the approximate policy function $\hat{U}(x,z)$. Then

$$V(x,z) = F(x, U^*(x,z), z) + \beta \int V\left(U^*(x,z) + \xi(z^1), z^1\right) Q(z, dz^1) \tag{55a}$$

$$W(x,z) = F(x, \hat{U}(x,z), z) + \beta \int W\left(\hat{U}(x,z) + \xi(z^1), z^1\right) Q(z, dz^1) \tag{55b}$$

Assume that with the numerical policy function $\hat{U}(x,z)$ the probability distribution over the state (x,z) converges weakly to the invariant measure $\hat{\mu}(x,z)$. We can interpret $\hat{\mu}(x,z)$ as the steady state distribution of (x,z), and the fact that $\hat{\mu}(x,z)$ is invariant under $Q(z,dz^1)$ implies that

$$\int f(x,z)\,\hat{\mu}(dx,dz) = \int\int f\left(\hat{U}(x,z) + \xi(z^1), z^1\right) Q(z,dz^1)\,\hat{\mu}(dx,dz) \tag{56}$$

for all measurable functions f. Subtracting (55b) from (55a), taking expectations, subtracting and adding $V(x,z)$ and applying (56) to V and W, the expected value loss can be written as

$$\int V(x,z) - W(x,z)\,\hat{\mu}(dx,dz) = \int \Big\{ F(x,U^*(x,z),z) - F(x,\hat{U}(x,z),z) + \beta\Big[\int \Big(V\left(U^*(x,z)+\xi(z^1),z^1\right) - V\left(\hat{U}(x,z)+\xi(z^1),z^1\right)\Big)\,Q(z,dz^1) + V(x,z) - W(x,z)\Big]\Big\}\,\hat{\mu}(dx,dz) \quad (57)$$

Using $\hat{U}(x,z) = U^*(x,z) - \delta(x,z)$ and applying a second order approximation to (57) gives

$$\begin{aligned}(1-\beta)&\int (V(x,z)-W(x,z))\,\hat{\mu}(dx,dz) \\ &= \int \Big[\delta(x,z)'F_2(x,u^*,z) - \frac{1}{2}\delta(x,z)'F_{22}(x,u^*,z)\delta(x,z) \\ +\beta&\int \left(\delta(x,z)'V_1\left(x^{*1},z^1\right) - \frac{1}{2}\delta(x,z)'V_{11}\left(x^{*1},z^1\right)\delta(x,z)\right)Q(z,dz^1)\Big]\,\hat{\mu}(dx,dz) + O\left({\eta_c}^3\right) \\ &= -\frac{1}{2}\int \delta(x,z)'M(x,z)\delta(x,z)\,\hat{\mu}(dx,dz) \\ &\quad - \delta(x,z)'\sum_{i=1}^{q} G_2^i(x,u^*,z)\lambda_i(x,z) + O\left({\eta_c}^3\right) \quad (58)\end{aligned}$$

where we have used the abbreviations $u^* = U^*(x,z)$, $x^{*1} = U^*(x,z) + \xi(z^1)$ and $M(x,z) \equiv F_{22}(x,u^*,z) + \beta\int V_{11}\left(x^{*1},z^1\right)Q(z,dz^1)$, as well as the first order condition $F_2(x,u^*,z) + \sum_i G_2^i(x,u^*,z)\lambda_i(x,z) + \beta\int V_1\left(x^{*1},z^1\right)Q(z,dz^1) = 0$. Observe that

$$\delta(x,z)'G_2^i(x,u^*,z)\lambda_i(x,z) \le 0 \quad (59)$$

because either $\lambda_i(x,z) = 0$, or $\lambda_i(x,z) > 0$ and $G^i(x,u^*,z) = 0$, and $\delta(x,z)'G_2^i(x,u^*,z)$ must be nonpositive since otherwise $\hat{U}(x,z) = u^* - \delta(x,z)$ would violate constraint i.

Now take as given an upper bound Δ^{ub} for $\mathrm{E}\left[V(x,z) - W(x,z)\right]$ that we have obtained by the method of Section 4. Assume we know that $M(x,z) + \alpha P$ is negative definite for all (x,z) in the support of $\hat{\mu}$, for a given positive definite matrix P and some $\alpha > 0$. To obtain such an α, we can either use an estimate of $\int V_{11}\,Q(z,dz^1)$ from the estimated value function, or use the fact (Maldonado and Svaiter 2001, slight generalization of their Lemma 3.1) that $V_{11} + \alpha P$ is negative definite if $F_{11} + \alpha P$ is. Combining (58) and (59) we get

$$\begin{aligned}&\mathrm{E}\left[\delta' P \delta\right] \le -\frac{1}{\alpha}\int \delta(x,z)'M(x,z)\delta(x,z)\,\hat{\mu}(dx,dz) \\ &\le \frac{2(1-\beta)}{\alpha}\int (V(x,z)-W(x,z))\,\hat{\mu}(dx,dz) + O\left({\eta_c}^3\right) \le \frac{2(1-\beta)\Delta^{ub}}{\alpha} + O\left({\eta_c}^3\right) \quad (60)\end{aligned}$$

Equ. (60) can be used to obtain variances and standard deviations of linear combinations of the policy errors.

5.2 Unbiased estimates of the policy error

To compute an unbiased estimate of the policy error, we start from Equ. (20) (for the sake of brevity, the following calculations only treat the case without inequality constraints)

$$R_{t+1} + F_{21,t}\delta_{t-1} + (F_{22,t} + \beta\,\mathrm{E}_t\,F_{11,t+1})\,\delta_t + \beta\,\mathrm{E}_t\,[F_{12,t+1}\delta_{t+1}] = O\left({\eta_c}^2\right) \tag{61}$$

Solving this for δ_t we obtain

$$\delta_t = -\left(F_{22,t} + \beta E_t[F_{11,t+1}]\right)^{-1}\left(F_{21,t}\delta_{t-1} + R_{t+1} + \beta\,\mathrm{E}_t\,[F_{12,t+1}\delta_{t+1}]\right) + O\left({\eta_c}^2\right) \tag{62}$$

Adopting again the approximation $\delta_T = 0$, we get

$$\delta_{T-1} = a^*_{T-1} + A^*_{T-1}\delta_{T-2} + O\left({\eta_c}^2\right) \tag{63}$$

where

$$a^*_{T-1} = -\left(F_{22,T-1} + \beta E_{T-1}[F_{11,T}]\right)^{-1} R_T \tag{64a}$$

$$A^*_{T-1} = -\left(F_{22,T-1} + \beta E_{T-1}[F_{11,T}]\right)^{-1} F_{21,T-1} \tag{64b}$$

This can be used as the starting point of the following recursion. Assume we are given the relationship

$$\delta_{t+1} = a^*_{t+1} + A^*_{t+1}\delta_t + O\left({\eta_c}^2\right) \tag{65}$$

where the vector a^*_{t+1} and the matrix A^*_{t+1} are known. Then it follows from (62) that

$$\delta_t = a^*_t + A^*_t\delta_{t-1} + O\left({\eta_c}^2\right) \tag{66}$$

where a^*_t and A^*_t are given by

$$a^*_t = -\left(F_{22,t} + \beta\,\mathrm{E}_t\left[F_{11,t+1} + F_{12,t+1}A^*_{t+1}\right]\right)^{-1}\left(R_{t+1} + \beta\,\mathrm{E}_t\left[F_{12,t+1}a^*_{t+1}\right]\right) \tag{67a}$$

$$A^*_t = -\left(F_{22,t} + \beta\,\mathrm{E}_t\left[F_{11,t+1} + F_{12,t+1}A^*_{t+1}\right]\right)^{-1} F_{21,t} \tag{67b}$$

Following this recursion we finally arrive at $\delta_0 = a^*_0$. Iterating Equs. (67) forwards, we prove the following proposition, which expresses the policy error δ_0 as an expectation of future Euler residuals:

Theorem 3.

$$\delta_0 = a^*_0 + O\left({\eta_c}^2\right) = \mathrm{E}_0 \sum_{s=0}^{T-1} C_s R_{s+1} + O\left({\eta_c}^2\right) \tag{68}$$

where

$$C_0 \equiv F_{12,0}^{-1} M_0 \tag{69}$$

$$C_{t+1} \equiv \beta C_t M_{t+1} \tag{70}$$

$$M_t \equiv -F_{12,t} \left(F_{22,t} + \beta \, \mathrm{E}_t \left[F_{11,t+1} + F_{12,t+1} A_{t+1}^* \right] \right)^{-1} \tag{71}$$

and the a_i^* and A_i^* are given by (64) and (67).

From (67b) we see that M_t satisfies the recursion

$$M_t = -F_{12,t} \left(F_{22,t} + \beta \, \mathrm{E}_t \left[F_{11,t+1} + M_{t+1} F_{21,t+1} \right] \right)^{-1} \tag{72}$$

If we can compute the matrices M_t, we can use (68) and Monte-Carlo techniques to compute the policy error δ_0 (or at least an unbiased estimate). The precision of the estimate can be arbitrarily increased by increasing the number of Monte-Carlo simulations, subject to an $O\left({\eta_c}^2\right)$-error from (61).

Comparing the recursions (67) with (39), we see that the difference lies in taking expectations as of time t. This is the reason why the estimate of δ_0 obtained here is unbiased, unlike its counterpart $\bar{\delta}_0$ of Section 4.2. This comes at a computational cost, however: while the recursions (39) can be computed along simulated paths, Equs. (67) are functional equations that have to be computed for the whole state space. Since the recursion (72) converges backwards in time, we can actually replace the M_t's in (68) by the solution to the time-invariant equation

$$\begin{aligned} M(x,z) = -F_{12}(x,z) \Big(& F_{22}(x,z) \\ & + \beta \int F_{11}(x + \xi(z^1), z^1) + M(x + \xi(z^1)) F_{21}(x + \xi(z^1), z^1) \, Q(z, dz^1) \Big)^{-1} \end{aligned} \tag{73}$$

where we write all the expressions as functions of the state (x, z). Equ. (73) is a nonlinear functional equation similar to a Bellman equation, but is considerably simpler because it does not involve any optimization. It can be obtained either recursively on a finite grid, or by projection methods.

6 Numerical examples

This section reports numerical results on a well known test problem, the one-dimensional stochastic neoclassical growth model. Numerical solutions to variants of this model have been intensively studied in the literature. I have chosen a one-dimensional model, since this can be solved at a finite grid with so high precision that it can serve as the "exact solution" for test purposes (Santos and Vigo-Aguiar 1998, Section 4).

6.1 One-dimensional stochastic growth

The social planner maximizes

$$\mathrm{E}\sum_{t=0}^{\infty}\beta^t\frac{c_t^{1-\gamma}-1}{1-\gamma} \tag{74}$$

subject to the constraints

$$c_t = \theta_t A k_t^{\alpha} - k_{t+1}$$
$$k_t \geq 0$$

where θ_t is an i.i.d. shock with log-uniform distribution

$$-\sigma \leq \log\theta_t \leq \sigma, \qquad \sigma > 0 \tag{75}$$

We bring this model into the form of Section 3 and solve it on a one-dimensional grid by defining[2]

$$x_t = \log(\theta_t A k_t^{\alpha}) \tag{76a}$$
$$u_t = \log(A k_{t+1}^{\alpha}) \tag{76b}$$
$$z_t = \log\theta_t \tag{76c}$$
$$\xi(z) = z \tag{76d}$$
$$F(x_t, u_t, z_t) = \frac{[e^{x_t} - (e^{u_t}/A)^{1/\alpha}]^{1-\gamma} - 1}{1-\gamma} \tag{76e}$$

The simulations use the parameter values $\beta = 0.95$, $\alpha = 0.4$, and $\sigma = 0.2$. I have chosen an extremely high variance of the shock, since we have seen in Section 4 that it is the uncertainty that raises problems in estimating the value and policy errors. A wide range of values for γ has been used, but always with $\gamma > 1$, which is necessary to make the function F concave in x and u (note that this restriction would not be necessary if we had chosen the representation given in Footnote 2). The constant A was set to $1/(\beta\alpha)$ so that the deterministic steady state capital stock is equal to 1.

The theory developed above is based on quadratic approximations, and so it might not be clear in practice whether the estimates are sufficiently accurate if there are strong nonlinearities with large third derivatives of the relevant functions. Since the stochastic growth model is known to be very well behaved, one might question the relevance of the results for less well-behaved problems. To deal with this problem without investigating more complicated models, I make the growth model as nonlinear as possible, first by

[2]Note that there are other ways to represent this model in our framework, the most straightforward representation is $x_t = k_t$, $u_t = k_{t+1}$, $z_t = \theta_t$, $\xi(z) = 0$ and $F(x_t, u_t, z_t) = z_t A x_t^{\alpha} - u_t$. I have chosen (76) since it reflects the fact that any combination of k_t and θ_t with the same $\theta_t A k_t^{\alpha}$) gives the same value to the household.

using very high degrees of risk aversion (up to $\gamma = 10$), and in one version also by introducing a kink in the marginal utility function at a value c^* close to the steady state level of consumption. More precisely, for $c \leq c^*$ the risk aversion parameter is some γ_0, for $c > c^*$ it is some γ_1, and at $c = c^*$ the utility function is scaled such that it is continuous and differentiable. The second derivative suddenly jumps downward at c^*. Note that this violates the assumptions made in Section 3; I have chosen this extreme specification in order to show that the assumptions on the smoothness of F will probably not be critical in practice. The switch point c^* was set at 1.02 times the deterministic steady state level of consumption.

Finally, in one version I have introduced an inequality constraint of the form $c_t \leq \phi e^{x_t}$ (note that this translates into a linear constraint in (x_t, x_{t+1})), where the constant ϕ was chosen so that, in the stationary distribution, the constraint is binding in about 65 percent of the periods. The constraint is similar to the familiar irreversibility of investment, but is somewhat artificial in the present model; its only purpose is to study the properties of the accuracy estimates in the case of inequality constraints.

6.2 Results

The "exact solution" was developed on a grid of 3137 points that conforms roughly to capital values $k \in (0.2, 5)$. I consider the following approximate solutions:

1. The solution obtained on discrete grids of 5, 25 and 97 points.

2. The linear approximation of u about the steady state. Since the state variable x is directly related to the log of the capital stock, this conforms to what is usually called a log-linear approximation. (In Tables 1–4, this approximation is denoted by "LinAppr".)

3. The policy $\hat{u}(x) = u^{exact}(x) + \phi$. Two different values of the constant ϕ were chosen, so as to make the size of the value function error close to that of the log-linear and the 25-point-grid approximation, respectively ("ConstErr1" and "ConstErr2").

4. The policy $\hat{u}(x) = u^{exact}(x) + \phi(x - x^{stst})$ where x^{stst} is the steady state value of x and ϕ was chosen to make the accuracy close to that of the 25-point-grid ("LinErr").

5. Same as the last approximation, but with an error only for values of x bigger than $0.95x^{stst}$. This creates a kink in the policy error ("Kink").

6. The policy $\hat{u}(x) = u^{exact}(x) + \phi x^{stst} \cdot \sin(\omega x)$. Here ω was chosen as 600π, making the error oscillate 3 times in a range of 1 percent of the capital stock.

Two different values of ϕ were chosen, analogously to "ConstErr" ("SinErr1" and "SinErr2").

7. A policy error proportional to $F_{11}(x)$. This could induce a bias in the estimated policy error (cf. the discussion at the beginning of Section 5). The size of the error was chosen as for "LinErr" (denoted by "Bias").

Items 1. and 2. of the list are approximations that arise from standard algorithms, while items 3. to 7. are artificially constructed in order to study the properties of the accuracy measure.

In the numerical calculations, the uniform distribution of θ was replaced by an 11-point distribution of equally likely points between $-\sigma$ and σ. The results are based on 1000 Simulations of 200 periods each. The shocks were not generated randomly, but as a sort of "subrandom" sequences, namely generalized Faure points.[3] This should increase the precision of the estimates. The "true value loss" was computed as the difference between the "exact value function" and the value obtained by an approximate solution, computed on a grid of 3137 points.

Tables 1 and 2 present results for the benchmark case of $\gamma = 5$. I have chosen the risk aversion parameter very big in the benchmark case, since for γ close to unity most of the approximate solutions are very accurate. The results in Tables 1 refer to average errors over the range of 0.5 to 2 times the steady state value of capital. The value loss Δ and its estimates are expressed in terms of proportional permanent changes in consumption (for example, a value loss of 0.005 is equivalent to a permanent 0.5 percent reduction in consumption).

The following conclusions emerge from the table. First, the estimates are very tight. The upper bound is never more than twice the true loss in value, and in most cases over-estimates by less than 20 percent. In two cases (Grid 5 and ConstErr1), the estimated upper bound is slightly less than the "true loss", even after accounting for the sampling error in computing both values. In both cases, the solution is relatively imprecise, so that the third order term in (35) is not yet negligible.

Second, the mean of the estimated policy error is close to the true error, so the estimate of $\bar{\delta}_0$ is not strongly biased in these examples. However, since the estimated error can be positive or negative, or close to zero by accident, it seems safer to look at the estimated *absolute* policy error. This gives, in all cases, a good indication of the size of the true policy error.

[3] Cf. Papageorgiou and Traub (1996). Thanks to Anargyros Papageorgiou and Joseph Traub for providing the FinDer software to generate the Faure points.

Third, estimating accuracy by means of Euler residuals is not equivalent to our method: policies with the same value loss can have very different maximum or average Euler residuals, compare for example "ConstErr2" and "LinErr".

Table 2 presents results for the same model, now not for averages over the state space, but starting from a specific value of the capital stock, namely 0.98 times the steady state. The same qualitative results hold. The "unbiased estimate" of Section 5 turns out to remove a significant part of the policy error bias in only two cases (LinErr, Bias), but slightly increases it in most other cases. I take this as an indication that the error from the quadratic approximation is more important than the bias resulting from the correlation between $\tilde{\delta}_{t+1}$ and $F_{12,t+1}$ etc. In any case, the estimated policy errors are so close that they do not provide good test cases to differentiate between the different error estimates.

Table 3 presents results for some other parameter values, including a version with kinked utility function. Again, the results are averages over the state space. We see that the qualitative results do not change when we vary the risk aversion parameter γ. Even in the case of the kink in the utility function, which violates the smoothness assumption made in Section 3, the method provides a tight upper bound. This suggests that the method is not very sensitive to the above smoothness assumptions.

Finally, Table 4 presents results for the model with the inequality constraint. The results show a very regular convergence behavior if the number of grid points of the approximate solution increases. While the statistic Δ^{ub} is actually below the true error for 5 grid points, and substantially overestimates the true error for 25 grid points, we obtain a tight upper bound for solutions on more grid points.

6.3 Computational cost

The computational cost of calculating the upper bound to the value function error is not trivial, but it is substantially lower than the cost of calculating the numerical solution itself, at least for medium or higher-dimensional problems. We only have to solve a series of linear-quadratic problems along realized paths, and calculate Euler residuals.

The recursive computation of the quadratic approximations is cheap and grows only cubically in the number of dimensions. More dimensions increase the computational effort mainly in two ways: it becomes more difficult to compute the Euler residuals, and we need more simulations to evaluate the expectation in (24). Both are problems of multidimensional integration. Computing (24) when x is univariate is the type of multidimensional integration that has been intensively studied in the literature on asset pricing (see e.g. Papageorgiou and Traub, 1996), and it has been found that Quasi-Monte-Carlo methods (for example generalized Faure points) performs best. To my

Approx.	Δ	Δ^{ub} (StDev)	Ave.Abs.Pol.Err	Est.Abs.Pol.Err. (StDev)	MaxER	AveER
Exact	n.a.	1.270e-10 (4.745e-12)	n.a.	2.964e-06 (7.833e-08)	6.368e-05	8.324e-06
Grid 5	2.187e-03	2.138e-03 (4.086e-05)	1.347e-02	1.254e-02 (1.199e-04)	4.348e-02	1.896e-02
Grid 25	5.909e-06	6.573e-06 (2.022e-07)	6.320e-04	6.637e-04 (1.333e-05)	9.431e-03	1.530e-03
Grid 97	7.287e-08	1.124e-07 (5.553e-09)	7.953e-05	9.649e-05 (2.727e-06)	1.732e-03	2.987e-04
LinQuad	5.472e-04	6.422e-04 (1.238e-05)	7.522e-03	7.478e-03 (1.409e-04)	9.285e-02	2.126e-02
ConstErr 1	3.674e-04	3.619e-04 (2.053e-06)	6.000e-03	5.841e-03 (6.031e-06)	1.801e-02	8.355e-03
ConstErr 2	5.678e-06	5.704e-06 (3.198e-08)	7.500e-04	7.531e-04 (7.192e-07)	2.227e-03	1.013e-03
LinErr	6.914e-06	7.387e-06 (4.092e-07)	1.347e-03	1.367e-03 (2.545e-05)	1.064e-02	2.974e-03
SinErr 1	4.200e-04	8.214e-04 (2.325e-05)	5.348e-03	7.327e-03 (1.673e-04)	6.943e-02	1.945e-02
SinErr 2	5.718e-06	1.102e-05 (2.994e-07)	6.192e-04	8.405e-04 (1.955e-05)	8.386e-03	2.260e-03
Kink	6.381e-06	7.070e-06 (5.340e-07)	9.212e-04	9.727e-04 (3.870e-05)	6.201e-03	1.607e-03
Bias	6.265e-06	6.495e-06 (4.579e-08)	7.009e-04	7.144e-04 (8.422e-06)	1.033e-03	7.818e-04

Δ is exact loss in value function, Δ^{ub} is estimated upper bound; StDev refers to sampling error

Ave.Abs.Pol.Err: average of absolute policy error $\delta_0(x)$ over state space

Est.Abs.Pol.Err: sample mean of absolute $\overline{\delta}_0(x)$

MaxER and AveER: maximal and average absolute Euler residuals over the state space

Table 1: Simulation results, growth model, $\gamma = 5$, $\sigma = 0.2$

Approx.	Δ	Δ^{ub} (StDev)	True Pol.Err	Est.Pol.Err. (StDev)	Est.Abs.Pol.Err.	Unbiased.Pol.Err.
Exact	n.a.	3.693e-10 (3.939e-12)	n.a.	-3.555e-06 (6.854e-08)	3.660e-06	-3.518e-06
Grid 5	6.565e-03	6.448e-03 (3.661e-05)	-1.725e-02	-1.579e-02 (5.318e-05)	1.579e-02	-1.521e-02
Grid 25	1.741e-05	1.890e-05 (1.930e-07)	-6.625e-04	-6.331e-04 (7.619e-06)	6.334e-04	-6.283e-04
Grid 97	1.874e-07	3.038e-07 (3.459e-09)	8.844e-05	8.608e-05 (1.947e-06)	9.135e-05	8.560e-05
LinQuad	1.639e-03	1.960e-03 (1.096e-05)	8.864e-03	9.618e-03 (3.607e-05)	9.618e-03	9.690e-03
ConstErr 1	1.080e-03	1.064e-03 (1.323e-06)	-6.000e-03	-5.847e-03 (5.233e-06)	5.847e-03	-5.674e-03
ConstErr 2	1.670e-05	1.677e-05 (2.034e-08)	-7.500e-04	-7.556e-04 (7.068e-07)	7.556e-04	-7.410e-04
LinErr	1.260e-05	1.424e-05 (2.087e-07)	7.741e-05	1.222e-05 (1.132e-05)	2.964e-04	4.764e-05
SinErr 1	1.178e-03	2.392e-03 (2.365e-05)	-1.428e-03	-1.241e-03 (1.810e-04)	4.699e-03	-1.152e-03
SinErr 2	1.615e-05	3.174e-05 (3.005e-07)	-1.653e-04	2.675e-04 (2.108e-05)	5.835e-04	2.666e-04
Kink	1.440e-05	1.636e-05 (3.940e-07)	1.161e-04	1.157e-04 (9.839e-06)	2.857e-04	1.396e-04
Bias	2.415e-07	2.695e-07 (4.375e-09)	-1.596e-05	-3.074e-05 (1.508e-06)	4.807e-05	-2.546e-05

Δ is exact loss in value function, Δ^{ub} is estimated upper bound

Est.Pol.Err: sample mean of estimated $\bar{\delta}_0(x_0)$

Est.Abs.Pol.Err: sample mean of absolute estimated $\bar{\delta}_0(x_0)$

Unbiased.Pol.Err: sample mean of estimated $\delta_0(x_0)$, formulas Section 5

Table 2: Simulation results, growth model, $\gamma = 5$, $\sigma = 0.2$; starting value $x_0 = 0.98$ times steady state

γ	Approx.	Δ	Δ^{ub} (StDev)	Ave.Abs.Pol.Err	Est.Pol.Err. (StDev)	MaxER	AveER
2	Exact	n.a.	1.052e-11 (1.090e-12)	n.a.	1.210e-06 (2.786e-08)	2.181e-05	6.706e-06
	Grid 5	3.669e-06	3.995e-06 (4.063e-07)	9.170e-04	9.412e-04 (1.876e-05)	1.004e-02	3.784e-03
	Grid 25	1.683e-07	2.067e-07 (2.342e-08)	2.388e-04	2.483e-04 (4.985e-06)	3.692e-03	1.222e-03
	LinAppr	4.146e-03	3.986e-03 (7.376e-04)	4.905e-02	4.838e-02 (9.924e-04)	8.265e-01	2.376e-01
10	Exact	n.a.	1.311e-09 (1.701e-11)	n.a.	6.320e-06 (2.044e-07)	1.270e-04	1.004e-05
	Grid 5	2.059e-02	1.654e-02 (9.545e-05)	2.988e-02	2.229e-02 (3.054e-04)	7.067e-02	2.007e-02
	Grid 25	4.683e-05	5.145e-05 (6.888e-07)	1.280e-03	1.313e-03 (3.763e-05)	7.342e-03	1.116e-03
	LinAppr	1.176e-02	1.544e-02 (1.263e-04)	2.643e-02	3.258e-02 (5.144e-04)	4.130e-01	3.697e-02
2, 10	Exact	n.a.	6.458e-11 (8.570e-12)	n.a.	2.098e-06 (5.690e-08)	6.555e-05	1.415e-05
	Grid 5	2.524e-02	2.348e-02 (1.571e-03)	4.612e-02	4.902e-02 (1.452e-03)	5.022e-01	2.290e-01
	Grid 25	3.210e-06	4.822e-06 (5.908e-07)	4.043e-04	5.291e-04 (1.654e-05)	1.635e-02	3.288e-03
	LinAppr	2.018e-03	2.789e-03 (6.147e-04)	2.899e-02	3.631e-02 (1.084e-03)	3.231e-01	1.246e-01

Notes: Two values of γ indicate kinked marginal utility function

Δ is exact loss in value function, Δ^{ub} is estimated upper bound; StDev refers to sampling error

Ave.Abs.Pol.Err: average of absolute policy error $\delta_0(x)$ over state space

Est.Abs.Pol.Err: sample mean of absolute $\bar{\delta}_0(x)$

MaxER and AveER: maximal and average absolute Euler residuals over the state space

Table 3: Simulation results, growth model, various levels of γ, $\sigma = 0.2$

Approx.	Δ	Δ^{ub} (StDev)	Ave.Abs.Pol.Err	Est.Abs.Pol.Err. (StDev)	MaxER	AveER	Constr.Err.
Exact	n.a.	1.218e-10 (5.262e-12)	n.a.	1.622e-06 (6.418e-08)	5.953e-05	8.130e-06	n.a.
Grid 5	6.925e-03	3.433e-03 (8.198e-05)	1.125e-02	9.726e-03 (2.216e-04)	1.102e+00	1.753e-01	1.023e-01
Grid 25	1.656e-06	4.130e-06 (1.877e-07)	2.191e-04	3.537e-04 (1.316e-05)	3.036e-02	4.272e-03	1.647e-01
Grid 49	5.801e-07	8.320e-07 (3.705e-08)	1.024e-04	1.451e-04 (5.563e-06)	9.137e-03	1.438e-03	7.992e-02
Grid 97	8.674e-08	9.684e-08 (4.412e-09)	3.215e-05	4.856e-05 (2.002e-06)	3.673e-03	4.817e-04	3.667e-02
Grid 193	2.915e-08	3.978e-08 (1.690e-09)	1.700e-05	2.878e-05 (1.185e-06)	1.159e-03	1.762e-04	1.564e-02
Grid 385	1.349e-08	1.543e-08 (6.896e-10)	1.064e-05	1.808e-05 (7.372e-07)	5.298e-04	1.138e-04	5.325e-03
Grid 769	3.010e-09	4.424e-09 (1.701e-10)	6.294e-06	1.009e-05 (4.019e-07)	2.521e-04	6.776e-05	3.500e-05

Δ is exact loss in value function, Δ^{ub} is estimated upper bound

Ave.Abs.Pol.Err: average of policy error δ_0 over state space

Est.Abs.Pol.Err: sample mean of absolute estimated $\bar{\delta}_0$

MaxER and AveER: maximal and average Euler residual over the state space

Constr.Err.: Frequency of constraint binding in exact solution and not in approximate solution, or vice versa

Table 4: Simulation results, growth model with inequality constraint, $\gamma = 5$, $\sigma = 0.2$

knowledge, there are no results that would tell us by how much the numerical effort goes up when x becomes multivariate. More numerical experiments have to be run to study this question. However, with Monte-Carlo and Quasi-Monte-Carlo methods the number of dimensions is usually not crucial.

Very likely, the main effect of dimensionality is on the computation of the Euler residuals. With up to about 6 dimensions, the most efficient way to compute Euler residuals is probably by quadrature. The computational burden will therefore depend essentially on the distribution of the stochastic shocks and the availability of quadrature grids. In an application with three dimensions or more, it may be very costly to compute the Euler residual at all points of the simulations. This can be avoided in many practical applications, since we can expect the integral term in the Euler residual to be a rather smooth function of the state variables. It may therefore be possible to find a parametric representation for it, in the spirit of the parameterized expectation algorithm of den Haan and Marcet (1990). Practically we would then compute the conditional expectation in the Euler residual at a certain number of points, and see whether it can be well explained as a smooth function of the state variables. This approximation can then be used in the simulations, which would drastically speed up the computations. Finally, one should remember that the theory did not depend on calculating the Euler residuals correctly: less precisely calculated residuals will simply result in a not so tight error bound, cf. Section 4.2. The effort to produce these estimates therefore depends again on the required precision on the upper bound.

The programming effort for any specific application is modest. The user only has to simulate the model, compute the Euler residuals and compute the Hessian of F at a given point of the state space. The latter task would be somewhat tedious to implement by hand. However, with algebraic software like Maple or with automatic differentiation software like ADIFOR, one can generate the required C- or Fortran code.

7 Conclusions: when to apply the method

The paper has derived estimates of the approximation error in the value as well as the policy function of a numerical solution to a dynamic optimization problem to which the exact solution is not known. Test applications have shown that both types of errors can be estimated with good precision.

It remains to discuss more generally the situations where the application of this method appears appropriate. Let me first point out that the method in its current form cannot be applied to models of discrete choice and to dynamic games. This restriction applies in fact to all evaluation methods based on Euler residuals. Apart from these

restrictions, the model is applicable to a wide variety of models analyzed in macroeconomics. To relate my method to other approaches, I briefly discuss four methods that I see as the main alternatives.

1. The statistic of Den Haan and Marcet (1994) is powerful in discriminating between different solution methods, and it indicates the dimensions along which a solution fails. It is also very easy to compute. However, it does not give an estimate of the size of the value or policy function error.

2. Estimating the maximal Euler residuals. As explained in Section 2, Euler residuals only give a one-period error, not an overall approximation error, since the effects of Euler residuals can add up over time. An evaluation based on average or maximal Euler residuals is therefore satisfactory only when the residuals, after suitable normalization, are trivially small, so that it is clear that the overall error is small enough (as, e.g., in Gaspar and Judd, 1997).

3. Concerning the error bounds of Santos (2000), which are based on the maximal Euler residuals, we already said that, by the nature of supremum arguments, they cannot be very tight. If the numerical solution is so precise that the bounds are sufficiently small for the purpose at hand, we can be confident about the solution; if not, it may be necessary to apply the present method to obtain tighter estimates.

4. Finally, one can estimate an upper bound on the value loss by a Bellman step on the estimated value function, compute the change ΔV, and use the well known formula $\left\|V - V^{exact}\right\| \leq \Delta V/(1-\beta)$. Maldonado and Svaiter (2001) show how to obtain an upper bound on the policy error from this information. They provide a result that is robust to any errors arising from approximating the value function. This method appears very natural if the model was solved by dynamic programming on a discrete grid. It then requires basically no further computations. Like Santos' method, it provides extremum bounds that are probably much looser than the estimates provided here. Since dynamic programming suffers from the curse of dimensionality, the approach is limited to problems of moderate dimension.

   One should note that it is possible to obtain an estimate of the value loss of an arbitrary numerical policy by computing the corresponding value function (this amounts to solving a linear functional equation) and applying a single Bellman step. However, apart from facing a curse-of-dimensionality problem, this step requires a discretization of the state space and the estimate of the value loss does not involve the error that arises from this discretization.

I therefore see the main application of my method in stochastic models with a medium- or high-dimensional state space, possibly with inequality constraints (an application to a standard imperfect market heterogenous agent model is provided in Reiter, 2001). In these situations, the dimensionality forbids exact computations on finite grids, and the inequality constraints induce kinks in the policy function and non-smooth value functions which make the application of smooth approximation methods problematic. These are also the models for which good error estimates are most urgently needed.

A Proofs of lemmas

Proof of Lemma 3. Premultiplying (48) by δ_t and taking expectations as of time t, we can write

$$\mathrm{E}_t\left[\delta_t' R_{t+1} + \delta_t' \hat{H}_{21,t}\bar{\delta}_{t-1} + \delta_t' \hat{H}_{22,t}\bar{\delta}_t + \beta\left(\delta_t' \hat{H}_{11,t+1}\bar{\delta}_t + \delta_t' \hat{H}_{12,t+1}\bar{\delta}_{t+1}\right)\right] = 0 \qquad (77)$$

Using (46) and the definition $\bar{\delta}_t = \delta_t - \tilde{\delta}_t$, multiplying (77) by β^t, summing over all time periods, and reordering as in (23), we get (recall that the function values $F_{1,t}$ etc. are always taken at the simulated paths $(\hat{x}_t, \hat{u}_t, z_t)$)

$$\mathrm{E}_0 \sum_{t=0}^{T} \beta^t \Big[\delta_{t-1}' F_{1,t} + \delta_t' F_{2,t} + \sum_i \left(\delta_{t-1}' G_{1,t}^i + \delta_t' G_{2,t}^i\right)\hat{\lambda}_t^i + Quad\left(\hat{H}, \delta, \delta; t\right) - Quad\left(\hat{H}, \delta, \tilde{\delta}; t\right)\Big] = O\left({\eta_c}^3\right) \qquad (78)$$

To analyze the term in $\hat{\lambda}_t$, consider the second order approximation

$$G^i\left(x_t^*, u_t^*, z_t\right) - G^i\left(\hat{x}_t, \hat{u}_t, z_t\right) = \delta_{t-1}' G_{1,t}^i + \delta_t' G_{2,t}^i + \frac{1}{2} Quad\left(G^i, \delta, \delta; t\right) + O\left({\eta_c}^3\right) \qquad (79)$$

Multiplying (79) by $\hat{\lambda}_t^i$ gives

$$\left(\delta_{t-1}' G_{1,t}^i + \delta_t' G_{2,t}^i\right)\hat{\lambda}_t^i = -\frac{1}{2}\hat{\lambda}_t^i Quad\left(G^i, \delta, \delta; t\right) + G^i\left(x_t^*, u_t^*, z_t\right)\hat{\lambda}_t^i + O\left({\eta_c}^3\right) \qquad (80)$$

since $G^i\left(\hat{x}_t, \hat{u}_t, z_t\right)\hat{\lambda}_t^i = 0$ by Assumption 1 ii). Inserting (80) into (78) and using $Quad\left(\hat{H}, \delta, \delta; t\right) = Quad\left(F, \delta, \delta; t\right) + \sum_{i=1}^{q} \hat{\lambda}_t^i Quad\left(G, \delta, \delta; t\right)$, we get

$$\mathrm{E}_0 \sum_{t=0}^{T} \beta^t \Big[\delta_{t-1}' F_{1,t} + \delta_t' F_{2,t} + Quad\left(F, \delta, \delta; t\right) + \frac{1}{2}\sum_{i=1}^{q} \hat{\lambda}_t^i Quad\left(G^i, \delta, \delta; t\right) + \sum_{i=1}^{q} G^i\left(x_t^*, u_t^*, z_t\right)\hat{\lambda}_t^i - Quad\left(\hat{H}, \delta, \tilde{\delta}; t\right)\Big] = O\left({\eta_c}^3\right) \qquad (81)$$

Subtracting (81) from (15), we obtain

$$\Delta^{appr} = \mathrm{E}_0 \sum_{t=0}^{T} \beta^t \left[-\frac{1}{2} Quad\left(\hat{H}, \delta, \delta; t\right) + Quad\left(\hat{H}, \delta, \tilde{\delta}; t\right) - \sum_{i=1}^{q} G^i\left(x_t^*, u_t^*, z_t\right)\hat{\lambda}_t^i\right] + O\left({\eta_c}^3\right) \qquad (82)$$

From the decomposition analogous to (26) we have

$$-\frac{1}{2}Quad\left(\hat{H},\delta,\delta;t\right)+Quad\left(\hat{H},\delta,\tilde{\delta};t\right)=-\frac{1}{2}\left[Quad\left(\hat{H},\bar{\delta},\bar{\delta};t\right)-Quad\left(\hat{H},\tilde{\delta},\tilde{\delta};t\right)\right] \quad (83)$$

Inserting (83) in (82), we obtain (50). □

In the following, we take explicitly into account that we are dealing with *sequences* of numerical approximations, indexed by n, where $n=1,2,\ldots$. We therefore write $\hat{x}_{t,n}$, $\hat{u}_{t,n}$ $\delta_{t,n}$, $\hat{\lambda}^i_n$ and η_n for the state variable, control variable, policy error, estimated Lagrange multiplier and supremum error of approximation n. In the assumptions of Lemmas 4 and 5, it should be understood that the constants Λ, $\bar{P}_\lambda$ and ϵ_0 hold uniformly for all n.

To proof Lemmas 4 and 5, we need

Lemma 6. *Under Assumption 1,* $G^i\left(x^*_t,u^*_t,z_t\right)\hat{\lambda}^i_t=O\left(\eta_n^2\right)$ *for all realizations of* z^T.

Proof. $\hat{\lambda}^i_n\left(\hat{x}_{t,n},z_t\right)>0$ implies that $G^i\left(\hat{x}_{t,n},\hat{u}_{t,n},z_t\right)=0$ and $G^i\left(x^*_t,u^*_t,z_t\right)=O\left(\eta_n^1\right)$ because of (9). Similarly, $G^i\left(x^*_t,u^*_t,z_t\right)>0$ implies that $\lambda^i\left(x^*_t,z_t\right)=0$ and $\hat{\lambda}^i_n\left(\hat{x}_{t,n},z_t\right)<\eta_n(1+\Lambda)$, for small enough η_n, where Λ is the Lipschitz constant from Assumption 1ii. Combining the two implications we get $G^i\left(x^*_t,u^*_t,z_t\right)\hat{\lambda}^i_n(\hat{x}_{t,n},z_t)=O\left(\eta_n^2\right)$. □

To obtain a result stronger than $O\left(\eta_n^2\right)$, note that, for any nonnegative random variable χ, we have

$$\mathrm{E}_0\left[\chi\right]=\mathrm{E}_0\left[\chi|\,\chi>0\right]\cdot Pr\left\{\chi>0\right\} \quad (84)$$

where $\mathrm{E}[\chi|A]$ denotes expectation of χ conditional on A. Note that
$\mathrm{E}_0\left[G^i\left(x^*_t,u^*_t,z_t\right)\hat{\lambda}^i_n(\hat{x}_{t,n},z_t)\middle|\,G^i\left(x^*_t,u^*_t,z_t\right)\hat{\lambda}^i_n(\hat{x}_{t,n},z_t)>0\right]=O\left(\eta_n^2\right)$ from Lemma 6.

Proof of Lemma 4. From Lemma 6 and (84), it suffices to show that the probability $Pr\left\{G^i\left(x^*_t,u^*_t,z_t\right)\hat{\lambda}^i(\hat{x}_t,z_t)>0\right\}$ goes to zero for $\eta_n\to 0$.

From (9) we have $\left\|G^i\left(\hat{x}_{t,n},\hat{u}_{t,n},z\right)-G^i\left(x^*_t,u^*_t,z\right)\right\|\leq\eta_n\cdot\overline{\nabla}$ for small enough η_n. Therefore

$$\begin{aligned}Pr\left\{G^i\left(x^*_t,u^*_t,z_t\right)>0\ \&\ \hat{\lambda}^i_n\left(\hat{x}_{t,n},z_t\right)>0\right\}&\leq Pr\left\{G^i\left(x^*_t,u^*_t,z_t\right)>0\ \&\ G^i\left(\hat{x}_{t,n},\hat{u}_{t,n},z_t\right)=0\right\}\\&\leq Pr\left\{G^i\left(x^*_t,u^*_t,z_t\right)\in\left(0,\eta_n\cdot\overline{\nabla}\right]\right\}\end{aligned}$$

Define $P_k\equiv Pr\left\{G^i\left(x^*_t,u^*_t,z_t\right)\in\left(\frac{1}{k+1},\frac{1}{k}\right]\right\}$. Then

$$Pr\left\{G^i\left(x^*_t,u^*_t,z_t\right)\in(0,1]\right\}=\sum_{i=1}^{\infty}P_i\equiv P^*\leq 1$$

Then

$$\lim_{k\to\infty}\sum_{i=k+1}^{\infty}P_i=P^*-\lim_{k\to\infty}\sum_{i=1}^{k}P_i=0$$

Define k_n as the largest k such that $1/k \geq \eta_n \cdot \overline{\nabla}$. Then, as $\eta_n \to 0$, it follows that $k_n \to \infty$ and therefore

$$Pr\left\{G^i\left(x_t^*, u_t^*, z_t\right) \in \left(0, \eta_n \cdot \overline{\nabla}\right]\right\} \leq \lim_{n\to\infty} \sum_{i=k_n}^{\infty} P_i \to 0 \tag{85}$$

□

Proof of Lemma 5. We know from the proof of Lemma 4 that

$$G^i\left(x_t^*, u_t^*, z_t\right) \hat{\lambda}_n^i(\hat{x}_{t,n}, z_t) > 0 \;\Rightarrow\; \hat{\lambda}_n^i\left(\hat{x}_{t,n}, z_t\right) \in (0, \eta_n(1+\Lambda)]$$

From the assumption in Lemma 5 it then follows that $Pr\left\{G^i\left(x_t^*, u_t^*, z_t\right) \hat{\lambda}^i(\hat{x}_t, z_t) > 0\right\} = O\left(\eta_n^1\right)$ for all t. The claim then follows from Lemma 6 and (84). □

B Discussion of Lemma 5

In this section, I want to show that the assumption underlying Lemma 5 holds under quite natural conditions.

Strictly speaking, one cannot provide more elementary conditions for (Equ. (53)) to hold, since the estimated Lagrange multipliers $\hat{\lambda}^i$ can be arbitrary functions, as long as they satisfy Assumption 1. To understand the logic of the assumption, it is therefore more instructive to look at the analogous condition for the exact multipliers:

$$Pr\left\{\lambda^i(x_t^*, z_t) \in (0, \epsilon]\right\} \leq \bar{P}_\lambda \epsilon \tag{86}$$

It appears plausible that a reasonable estimate $\hat{\lambda}^i$ will satisfy (53) if (86) is fulfilled.

We will first reduce (86) to a similar condition on the policy functions. For this it is necessary to introduce the concept of an unrestricted policy function. We know that $U^*\left(x, z_0\right)$ solves the HJB equation

$$V(x,z) = \max_{u: G(x,u,z)\geq 0} \left\{F(x,u,z) + \beta \int V\left(u + \xi(z^1), z^1\right) Q(z, dz^1)\right\} \tag{87}$$

Now define $U_{-i}^*\left(x, z\right)$ as the solution to

$$\max_{u: G^j(x,u,z)\geq 0, j\neq i} \left\{F(x,u,z) + \beta \int V\left(u + \xi(z^1), z^1\right) Q(z, dz^1)\right\} \tag{88}$$

where the maximization is over the u that satisfy all restrictions except i. This definition conforms to the problem where in the present period, the constraint i is ignored, while in all the future periods the constraint is respected (so that the same value function is used in (87) and in (88).

We can now establish the following lemma:

Lemma 7. *Assume:*

i) The function $F(x,u,z)$ is strongly concave in u in the sense that there exists an $\alpha > 0$ such that $F(x,u,z) + \alpha u'u$ is concave in (x,u) for all z.

ii) There exist constants $\epsilon_0 > 0$ and $\bar{\pi}$, s.t. for all u, i, t and $\epsilon < \epsilon_0$,

$$Pr\left\{ \left\| U^*\left(u + \xi(z_t), z_t\right) - U^*_{-i}\left(u + \xi(z_t), z_t\right) \right\| \in (0, \epsilon] \;\middle|\; z_{t-1} \right\} < \bar{\pi}\epsilon$$

Then Equ. (86) holds.

Proof. Step 1: First we prove a relationship between the value of the Lagrange multiplier and the distance of the optimal restricted policy from the unrestricted policy.

Define $W(x,u,z) \equiv F(x,u,z) + \beta \int V\left(u + \xi(z^1), z^1\right) Q(z, dz^1)$. Since V is concave, Assumption i) of the lemma implies that $W + \alpha u'u$ is concave.

Take any $(x_0, z_0) \in \mathcal{Y}$, and set $u^* \equiv U^*(x_0, z_0)$ and $\tilde{u} \equiv U^*_{-i}(x_0, z_0)$. Define $\delta_u \equiv u^* - \tilde{u}$. Since $\tilde{u}$ solves (88), the directional derivative $W_2(x_0, \tilde{u}, z_0)'\delta_u$ must be equal to zero, because the control u at $(x_0, \tilde{u}, z_0)$ is not restricted in the direction δ_u. Towards u^*, this follows from the fact that the set of feasible controls is convex. In the opposite direction, it cannot be restricted because, again for convexity, u^* would not satisfy this restriction, which is impossible.

Now define the scalar function $f(t) \equiv W_2(x_0, \tilde{u} - t\delta_u, z_0)'\delta_u$. We have just shown that $f(0) = 0$ and then we get

$$\begin{aligned} W_2(x_0, u^*, z_0)\delta_u = f(1) = \int_0^1 \frac{df(t)}{dt}\, dt = -\int_0^1 \delta_u' W_{uu}(x_0, \tilde{u} - t\delta_u, z_0)\delta_u \, dt \\ \geq \alpha \left\| \delta_u \right\|^2 \end{aligned}$$

Now assume that the constraint i is relaxed to $G^i(x,u,z_0) \geq -\epsilon$ with $\epsilon > 0$. Then the policy $u^* + \gamma\delta_u$ is feasible where

$$\gamma = \frac{\epsilon}{\overline{\nabla} \left\| \delta_u \right\|} \leq \frac{\epsilon}{\left\| \delta_u \right\| \cdot \left\| G_2(x_0, u^*, z_0) \right\|}$$

Then the change in the value function which is possible from the relaxation of the constraint is at least $W(x_0, u^* + \gamma\delta_u, z_0) - W(x_0, u^*, z_0) = \gamma W_2(x_0, u^*, z_0)'\delta_u$, where the equality is up to a first order approximation in ϵ. If we define $u^*(\epsilon)$ as the optimal u under the relaxed constraint, so that $V(x_0, z_0) = W(x_0, u^*(0), z_0)$, we obtain

$$\lambda^i(x_0, z_0) = \left. \frac{dW(x_0, u^*(\epsilon), z_0)}{d\epsilon} \right|_{\epsilon=0} \geq \gamma W_2(x_0, u^*, z_0)'\delta_u \geq \alpha \left\| \delta_u \right\| / \overline{\nabla} \tag{89}$$

Step 2: From Step 1 we have that

$$Pr\left\{ \lambda^i(x_0, z_0) \in (0, \epsilon] \right\} \leq Pr\left\{ \left\| U^*(x_0, z_0) - U^*_{-i}(x_0, z_0) \right\| \in \left(0, \epsilon\overline{\nabla}/\alpha\right] \right\}$$

Since this is true for any (x_0, z_0), we can substitute $u + \xi(z_t)$ for x_0 and z_t for z_0, and then the assumptions of the lemma imply that the latter probability (conditional on z_{t-1}, and therefore also the unconditional probability, by the law of iterated expectations) is smaller than $\bar{\pi}\overline{\nabla}\epsilon/\alpha$. This shows that (86) holds with constant $\bar{P}_\lambda = \bar{\pi}\overline{\nabla}/\alpha$. □

Assumption ii) of Lemma 7 is critical. Whether it holds or not depends mainly on whether the unrestricted policy function and the slope of the constraint intersect with a nonzero angle or not. Let us illustrate this with an example of one state variable x, one control u and one constraint $G^1(x, u) \geq 0$, where $x_t = u_{t-1} + \xi_t$. Assume there is exactly one x° such that $G(x^\circ, U^*_{-1}(x^\circ)) = 0$. Assume further that $\frac{d\, U^*_{-1}(x^\circ)}{d\, x}$ differs from $-\frac{G_x(x^\circ, U^*(x^\circ))}{G_u(x^\circ, U^*(x^\circ))}$, the slope of the constraint at x°, by at least $\phi > 0$. Then $\|U^*(x) - U^*_{-1}(x)\| \in (0, \epsilon]$ implies $\|x - x^\circ\| < 2\epsilon/\phi$ for small enough ϵ. If ξ_t is a continuous random variable with density function bounded by $\bar{p}$, this implies that $Pr\left\{\|U^*(u + \xi_t) - U^*_{-1}(u + \xi_t)\| \in (0, \epsilon]\right\} < 4\bar{p}\epsilon/\phi$, verifying Assumption ii). In contrast, if $\frac{d\, U^*_{-1}(x^\circ)}{d\, x} = -\frac{G_x(x^\circ, U^*(x^\circ))}{G_u(x^\circ, U^*(x^\circ))}$, it is straightforward to construct examples where Equ. (86) does not hold.

The ideas of this example can be carried over to higher dimensions by requiring that the graph of the unrestricted policy function and the constraint surface intersect with different slopes, in a direction of the state space where the random shocks generate enough variability. Since a precise and reasonably general formulation of these ideas is rather tedious and requires a lot of notation, I leave it with the above example, which should illustrate the relevant issues.

References

Benitez-Silva, H., Hall, G., Hitsch, G. J., Pauletto, G. and Rust, J.: 2000, A comparison of discrete and parametric approximation methods for continuous-state dynamic programming problems, mimeo.

Billingsley, P.: 1986, *Probability and Measure*, 2nd edn, John Wiley & Sons, New York.

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

Collard, F., Fève, P. and Juillard, M.: 2000, Simulation of stochastic non-linear models: Testing the accuracy of approximation, Cepremap.

den Haan, W. J. and Marcet, A.: 1990, Solving the stochastic growth model by parameterizing expectations, *Journal of Business and Economic Statistics* **8**(1), 31–34.

Den Haan, W. J. and Marcet, A.: 1994, Accuracy in simulations, *Review of Economic Studies* **61**(1), 3–17.

Gaspar, J. and Judd, K. L.: 1997, Solving large-scale rational-expectations models, *Macroeconomic Dynamics* **1**(1), 45–75.

Judd, K. L.: 1992, Projection methods for solving aggregate growth models, *Journal of Economic Theory* **58**(2), 410–52.

Maldonado, W. L. and Svaiter, B.: 2001, On the accuracy of the estimated policy function using the Bellman contraction method, mimeo.

Papageorgiou, A. and Traub, J.: 1996, Beating Monte Carlo, *Risk* **9**(6), 63–5.

Reiter, M.: 2001, Recursive solution of heterogenous agent models, mimeo, Universitat Pompeu Fabra.

Santos, M. S.: 2000, Accuracy of numerical solutions using the Euler equation residuals, *Econometrica* **68**(6), 1377–1402.

Santos, M. S. and Vigo-Aguiar, J.: 1998, Analysis of a numerical dynamic programming algorithm applied to economic models, *Econometrica* **66**(2), 409–26.

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

Taylor, J. B. and Uhlig, H.: 1990, Solving nonlinear stochastic growth models: A comparison of alternative solution methods, *Journal of Business and Economic Statistics* **8**(1), 1–17.