

Using Nonlinear Model Predictive Control for Dynamic Decision Problems in Economics*

January 20, 2015

Lars Grüne,¹ Willi Semmler² and Marleen Stieler³

Abstract: This paper presents a new approach to solve dynamic decision models in economics. The proposed procedure, called Nonlinear Model Predictive Control (NMPC), relies on the iterative solution of optimal control problems on finite time horizons and is well established in engineering applications for stabilization and tracking problems. Only quite recently, extensions to more general optimal control problems including those appearing in economic applications have been investigated. Like Dynamic Programming (DP), NMPC does not rely on linearization techniques but uses the full nonlinear model and in this sense provides a global solution to the problem. However, unlike DP, NMPC only computes one optimal trajectory at a time, thus avoids to grid the state space and for this reason the computational demand grows much more moderately with the space dimension than for DP. In this paper we explain the basic idea of NMPC, give a proof concerning the accuracy of NMPC for discounted optimal control problems, present implementational details, and demonstrate the ability of NMPC to solve dynamic decision problems in economics by solving low and high dimensional examples, including models with multiple equilibria, tracking and stochastic problems.

Keywords: Complex decision models, long and short horizon models, dynamic optimization, multiple equilibria, regime changes, NMPC

*Willi Semmler would like to thank the Fulbright Commission for research support. We also want to thank Vladimir Veliov and his research group at the University of Technology, Vienna. Thanks go to Chris Sims for very helpful discussion. Versions of the paper have been presented at the 6th and 7th international conferences on Computation in Finance and Economics, Oviedo, Spain, and London, UK. Comments from the participants and Michel Juillard are gratefully appreciated. Lars Grüne was supported by the European Union under the 7th Framework Program FP7–PEOPLE–2010–ITN, Grant agreement nr. 264735–SADCO. Marleen Stieler was supported by the International Doctorate Program “Identification, Optimization and Control with Applications in Modern Technologies” within the Elite Network of Bavaria.

¹Mathematical Institute, University of Bayreuth, Germany, email: lars.gruene@uni-bayreuth.de

²New School For Social Research, New York, University of Bielefeld, and Research Associate at the Center for European Economic Research, Mannheim, email: semmlerw@newschool.edu

³Mathematical Institute, University of Bayreuth, Germany, email: marleen.stieler@uni-bayreuth.de

1 Introduction

The lack of closed form solutions of dynamic decision models with optimizing agents has generated a large number of computational methods to solve such models. A detailed discussion of a variety of numerical methods and accuracy tests are provided in Santos and Vigo-Aguiar (1998), Judd (1998), Juillard and Villemot (2011) and Grüne and Semmler (2004). The latter have proposed Dynamic Programming (DP), with grid refinement, cf. Grüne (1997), to solve a family of continuous and discrete time dynamic models with optimizing agents. DP provides the value function and the control variable in feedback form, even for rather complex problems.

In DP a global solution to the optimal control problem is found by first computing an approximation to the optimal value V and then computing the optimal control from V , see Grüne and Semmler (2004). Yet, since DP computes the value and policy function at each point of a grid of the state space, it has the disadvantage that even with an adaptive choice of the grid its numerical effort typically grows exponentially with the dimension of the state variable. Hence, already for moderate state dimensions it may be impossible to compute a solution with reasonable accuracy.

This paper illustrates how Nonlinear Model Predictive Control (NMPC) can be used as an alternative approach to solve dynamic decision models in economics. NMPC is a well known method in control engineering which is frequently used in industrial practice, particularly in chemical process engineering. Traditionally, NMPC is applied to optimal feedback stabilization problems, see, e.g., Rawlings and Mayne (2009) or Grüne and Pannek (2011) and the references therein. Recently, however, the application of NMPC to more general optimal control problems has attracted considerable attention, see, e.g., Amrit et al. (2011); Angeli et al. (2009); Angeli and Rawlings (2010); Diehl et al. (2011); Grüne (2013); Grüne and Stieler (2014) for undiscounted optimal control problems. Similar to DP, NMPC can solve nonlinear dynamic decision problems globally without having to resort to local approximations by linearization techniques.

However, unlike DP the solution is not found on a grid in state space. Rather, an infinite horizon trajectory is synthesized by putting together pieces of finite horizon optimal trajectories, which implies that the numerical effort of the approach scales much more moderately with the state dimension. This approach, termed *receding horizon control* in control engineering, is in fact not unknown in economics. In the economic literature, it is known as *sliding or rolling planning*, see, e.g., Kaganovich (1985) and the references therein. However, in the economic context we are only aware of applications of this approach to linear models. The contribution of this paper is to demonstrate that NMPC also applies to nonlinear problems in dynamic decision making in economics. To this end, we establish a convergence result for discounted optimal control problems and illustrate its performance by applying it to several economic decision models.

Assuming that a reliable numerical solver for finite horizon optimal control problems is available⁴, the main source of errors in NMPC is the difference between the optimal trajectories of finite and infinite horizon optimal control problems. In Section 3, we show that for discounted optimal control problems with small discount factor and for problems satisfying

⁴For a discussion of this aspect see Section 7.

the so called turnpike property, cf. McKenzie (1986), for sufficiently long finite horizons, NMPC yields approximately infinite horizon optimal trajectories. Unlike other numerical errors, like, e.g., interpolation errors in DP, this source of errors allows for a precise economic interpretation. In fact, the mismatch between the true solution of an infinite horizon decision problem and its NMPC solution is due to the fact that the decision for the control to be implemented at the next time step is taken by looking at the problem on a truncated time horizon, i.e., with a particular form of decision making under limited information.

Sims (2005, 2006), in a series of research papers, showed that agents make decisions under limited information: The information is either not available or the agents respond imprecisely to the available information. In this context, we can interpret the gap between the infinite horizon solution and the NMPC solution as induced by the agents' decision making using only limited information. As such, the abstract convergence results from Section 3 have a self-evident economic interpretation: if the agents information and information processing capacity increases this is likely to approximate better the infinite horizon decision making⁵, which is reflected, e.g., in the examples in Sections 4.1 and 5.1 below. One might even use NMPC to systematically study the effects of decision making for this particular form of rational inattention.

Though we do not elaborate further on the latter aspect in this paper, we would like to make some remarks on the finiteness of the decision horizon. The argument could be made that if the agents come close to the final period they will sell all their assets which will impact the last period's outcome. Yet, the way the NMPC solution procedure is set up, only the first decision step is implemented. If the decision horizon is N , then one is, in the closed loop solution, always $N - 1$ periods away from the final decision. Hence, one never sees the effects which appear at the end of the decision horizon. This can be formalized using the turnpike property, which then allows to prove a formal convergence result, see Section 3. Under the appropriate conditions, this property holds without using the salvage value of the finite horizon model. If, however, the decision horizon is short it might be beneficial to take into account the salvage value, provided it can be determined in a reasonable way. Likewise, information about optimal steady states may be incorporated into the NMPC algorithm via terminal constraints which may be useful for short decision horizons. Yet, as the decision horizon becomes larger, typically there is no need for taking the salvage value or information about steady state into account since NMPC already approximates well the infinite horizon decision model. One of the issues will thus be how large the decision horizon N needs to be, see Sections 3 as well as the discussion for the examples in Sections 4.1 and 5.1.

In this paper, we evaluate the performance of NMPC analytically and via computer simulations for a selection of dynamic decision models in economics. Particularly, we extend the economic MPC results from the literature by considering discounted optimal control problems, both analytically and by studying a number of examples by means of numerical simulations. In order to study the accuracy of NMPC for approximating discounted infinite horizon problems, we first want to test our algorithm by studying the well-known basic growth model of Brock and Mirman (1972) type, for which the exact solution is known,

⁵Sims notes "... the capacity-constrained agent's behavior approximates that of a fully optimizing agent, but with a tight capacity constraint his behavior will be much more weakly correlated with external information than the behavior of a fully optimizing agent would be." (Sims; 2006, p. 158)

and a recent DSGE extension of it. To study the Brock et al model allows us to judge the accuracy of our numerical method for a model with short decision horizon, and to explore what the new method can contribute.

As mentioned above, there are, in the economic literature, more complicated dynamic models with optimizing agents which have been a challenge to commonly used numerical techniques. These are models with multiple equilibria, regime changes in the dynamics, models for tracking the dynamics of the state variables over a finite horizon and dynamic models with parameter uncertainties and learning⁶. Examples of such challenges to numerical solution methods can be found in the literature on the dynamic decision problems of the firm,⁷ economic growth,⁸ macrodynamics of DSGE type,⁹ in models with exhaustible resources and in ecological management problems.¹⁰ Our paper here studies some of the proto-type models from some of those areas and applies the proposed NMPC to find the global solution.

The remainder of the paper is organized as follows. Section 2 describes the basic strategy of NMPC. In Section 3 we provide convergence results for NMPC applied to infinite horizon discounted optimal control problems. Section 4 solves one dimensional control problems with one and two decision variables. Here we study the basic growth model for which the exact solution is known, so that we address the accuracy problem. We also study a DSGE extension of it. In Section 5 we then study two dimensional dynamic optimization models, specifically a model with two domains of attraction and a threshold, and another one where the paths of the state variables need to be tracked. In Section 6 we apply NMPC to a higher dimensional problem, by discussing a five dimensional integrated assessment model of growth and climate change. Section 7 gives details of the numerical implementation and discusses some known pitfalls of the NMPC method. Section 8 provides an outlook how NMPC can be used for stochastic problems, and Section 9 concludes the paper.

2 Nonlinear model predictive control

In this section we describe the basic principles of the NMPC method. Further implementational details are discussed in Section 6.

2.1 Problem formulation

We consider infinite horizon discounted optimal control problems, either given in continuous time $t \in \mathbb{R}_0^+$ by

$$V(x_0) := \max_{u \in \mathcal{U}} \int_0^\infty e^{-\delta t} g(x(t), u(t)) dt \quad (2.1)$$

⁶For the latter see Bréchet et al. (2012)

⁷See Feichtinger et al. (2001) and Haunschmied et al. (2003).

⁸In the latter type of models a convex-concave production function arises which leads to thresholds separating paths to low per capita income (poor) countries and high per capita income (rich) countries, see Skiba (1978) and Azariadis and Drazen (1990).

⁹See Farmer et al. (2009)

¹⁰See in particular Brock and Starrett (1999)

where

$$\frac{d}{dt}x(t) = f(x(t), u(t)), \quad x(0) = x_0 \in X \subseteq \mathbb{R}^n \quad (2.2)$$

or in discrete time $t \in \mathbb{N}_0$ given by

$$V(x_0) := \max_{u \in \mathcal{U}_d} \sum_{k=0}^{\infty} \beta^k g(x(k), u(k)) \quad (2.3)$$

where

$$x(k+1) = \varphi(x(k), u(k)), \quad x(0) = x_0 \in X \subseteq \mathbb{R}^n \quad (2.4)$$

and \mathcal{U} and \mathcal{U}_d are appropriate sets of control functions and control sequences, respectively, and $X \subseteq \mathbb{R}^n$ is the part of the state space we are interested in.

NMPC as described in the sequel applies to discrete time problems, hence the continuous time problem needs to be discretized in time in order to apply the method¹¹. To this end, the continuous time optimal control problem (2.1)–(2.2) is replaced by a first order discrete time approximation given by

$$V_h(x_0) := \max_{u \in \mathcal{U}_d} J_h(x_0, u), \quad J_h(x, u) := \sum_{k=0}^{\infty} \beta^k g_h(\tilde{x}(k), u(k)) \quad (2.5)$$

where $\beta = e^{-\delta h}$, $g_h(x, u) = hg(x, u)$ and $\tilde{x}(k)$ is defined by the discrete dynamics

$$\tilde{x}(k+1) = \varphi_h(\tilde{x}(k), u(k)), \quad \tilde{x}(0) = x_0, \quad (2.6)$$

where $h > 0$ is the discretization time step and φ_h is a numerical approximation to the continuous time solution of (2.2) at time h . If the original problem is of type (2.3)–(2.4), then it is already in the form (2.5)–(2.6) with $h = 1$ and $\tilde{x}(k) = x(k)$. Since in the remainder of this section we exclusively deal with discrete time problems, in order to simplify the notation we will omit the indices h and d and the tilde on $\tilde{x}(k)$. In exchange, we will use indices ∞ and N in order to indicate optimization on infinite and finite horizons, respectively.

2.2 The idea of NMPC

The idea of NMPC now lies in replacing the maximization of the discrete time infinite horizon functional

$$\max_{u \in \mathcal{U}} J_{\infty}(x(0), u) \quad \text{where } J_{\infty}(x(0), u) := \sum_{k=0}^{\infty} \beta^k g(x(k), u(k)) \quad (2.7)$$

from (2.3) by the iterative maximization of finite horizon functionals

$$\max_{u \in \mathcal{U}} J_N(x(0), u) \quad \text{where } J_N(x(0), u) := \sum_{k=0}^{N-1} \beta^k g(x(k), u(k)) \quad (2.8)$$

¹¹This approach is similar to the first step of the semi-Lagrangian discretization technique for the DP method going back to Capuzzo Dolcetta (1983) and Falcone (1987) and also described in Grüne and Semmler (2004).

for a truncated finite horizon $N \in \mathbb{N}$ with given initial value $x(0) \in X \subseteq \mathbb{R}^n$ and x_i generated by the usual dynamics $x(k+1) = \varphi(x(k), u(k))$ for $k = 0, 1, 2, \dots, N-1$. Note that in the standard case, neither β nor g nor φ changes when passing from (2.7) to (2.8), only the optimization horizon is truncated, though changes can be allowed for in extensions, see the discussion in the second past paragraph of Section 3.3 and the Example in Section 4.2, below.

Problems of type (2.8) can be efficiently solved numerically by converting them into a static nonlinear program (NLP) and solving them by NLP solvers, implementation details are discussed in Section 7, below.

Given an initial value $x_N(0)$, NMPC now generates solutions $x_N(i)$ on an infinite time horizon by iteratively solving (2.8) as follows:

- (1) for $i = 0, 1, 2, 3, \dots$
- (2) solve (2.8) with initial value $x(0) := x_N(i)$ and denote the resulting optimal control sequence by $u_{N,i}^*(\cdot)$
- (3) set $u_N(i) := u_{N,i}^*(0)$ and $x_N(i+1) := \varphi(x_N(i), u_N(i))$
- (4) end of for-loop

This algorithm yields an infinite trajectory $x_N(i)$, $i = 0, 1, 2, 3, \dots$ whose control sequence $u_N(i)$ consists of all the first elements $u_{N,i}^*(0)$ of the optimal control sequences for the finite horizon subproblems (2.8). In what follows, we refer to the finite horizon optimal trajectories corresponding to $u_{N,i}^*(\cdot)$ computed in Step (2) as the *open loop* trajectories while the trajectory $x_N(i)$ computed in Step (3) will be referred to as the *closed loop* trajectory.

3 Convergence analysis

Results which show that the value of an infinite horizon performance criterion evaluated along an NMPC solution approximates the infinite horizon optimal value for sufficiently large horizon have been obtained for various types of undiscounted optimal control problems. We refer to, e.g., Amrit et al. (2011); Angeli et al. (2009); Angeli and Rawlings (2010); Diehl et al. (2011) for averaged performance criteria, Grüne (2013); Grüne and Stieler (2014) for non-averaged criteria and Sections 5.4 and 6.5 of Grüne and Pannek (2011) for stabilizing problems. In this section we investigate the infinite horizon performance of the trajectories generated by the NMPC trajectories for discounted optimal control problems. To this end, we consider the discounted infinite horizon functional

$$J_\infty^{cl}(x, \mu_N) := \sum_{k=0}^{\infty} \beta^k g(x_{\mu_N}(k, x), \mu_N(x_{\mu_N}(k, x)))$$

evaluated along the closed loop MPC trajectory. Our goal is to show that this value is close to the optimal value

$$V_\infty(x) := \max_{u \in \mathcal{U}} J_\infty(x, u)$$

under suitable conditions and that the difference between $J_\infty^{cl}(x, \mu_N)$ and $V_\infty(x)$ converges to 0 as $N \rightarrow \infty$. For this analysis, we will also utilize the finite horizon optimal value

function

$$V_N(x) := \max_{u \in \mathcal{U}} J_N(x, u).$$

An important tool in the sequel is the dynamic programming principle, which states that

$$V_N(x) = J_K(x, u_N^*) + \beta^K V_{N-K}(x_{u_N^*}(K, x))$$

in the finite horizon case for all $K \leq N$ and

$$V_\infty(x) = J_K(x, u_\infty^*) + \beta^K V_\infty(x_{u_\infty^*}(K, x))$$

in the infinite horizon case for all $K \in \mathbb{N}$. Here u_N^* and u_∞^* denote the finite and infinite horizon optimal control sequences for initial value x .

3.1 The case of small discount factors

If the discount factor $\beta \in (0, 1)$ is small, i.e., not too close to 1, then it is rather easy to establish a convergence estimate for MPC. Assuming that V_N and V_∞ are bounded by a constant M on X , a little computation using the dynamic programming principle reveals that

$$|V_N(x) - V_\infty(x)| \leq \beta^N M$$

for all $x \in X$. Using the dynamic programming principle once more then yields

$$g(x, \mu_N(x)) = V_N(x) - \beta V_{N-1}(f(x, \mu_N(x))) = V_\infty(x) - \beta V_\infty(f(x, \mu_N(x))) + R$$

where the remainder term R satisfies $|R| \leq 2\beta^N M$. From this we obtain

$$\begin{aligned} J_\infty^{cl}(x, \mu_N) &= \sum_{k=0}^{\infty} \beta^k g(x_{\mu_N}(k, x), \mu_N(x_{\mu_N}(k, x))) \\ &\leq \sum_{k=0}^{\infty} \beta^k \left(V_\infty(x_{\mu_N}(k, x)) - \beta V_\infty(x_{\mu_N}(k+1, x)) + 2\beta^N M \right) \\ &= V_\infty(x) + \frac{2\beta^N M}{1 - \beta} \end{aligned}$$

due to the boundedness of V_∞ . Since $2\beta^N \rightarrow 0$ as $N \rightarrow \infty$, this shows the desired estimate.

3.2 The case of discount factors close to 1

While the reasoning of the previous section also applies if β is close to 1, the convergence of $\beta^N \rightarrow 0$ will be very slow and the resulting estimate will be overly pessimistic. For instance, in the example in Section 4.1 for $N = 10$ we obtain an accuracy of about 10^{-8} . However, since $\beta = 0.95$ we have $\beta^{10} \approx 0.6$, i.e., the factor in the error term is by many orders of magnitude larger than the actual error of NMPC and does by no means explain its almost perfectly optimal performance.

For this reason, in this section we provide an alternative convergence proof which relies on the turnpike property rather than on the size of the discount factor. For stating the

respective assumption, we need to introduce some notation. As above, u_∞^* and u_N^* denote the infinite and finite horizon optimal control sequences, respectively. A point $x^* \in X$ is called (infinite horizon) *optimal equilibrium*, if $x_{u_\infty^*}(k, x^*) = x^*$ for all $k \in \mathbb{N}$. A function $\alpha : \mathbb{R}_0^+ \rightarrow \mathbb{R}_0^+$ is of class \mathcal{K}_∞ if it is continuous, unbounded and strictly increasing with $\alpha(0) = 0$. Moreover, as usual we denote the open ball around $x \in X$ with radius $\varepsilon > 0$ as $\mathcal{B}_\varepsilon(x)$.

Now we make the following assumptions:

Assumption 3.1 There is an optimal equilibrium $x^* \in X$ with corresponding control value $u^* \in U$ satisfying $f(x^*, u^*) = x^*$ and the following properties.

- (i) x^* is asymptotically stable for the infinite horizon problem in the following sense: there exists $\rho \in \mathcal{L}$ such that the optimal trajectories $x_{u_\infty^*}(k, x)$, $x \in X$ satisfy

$$\|x_{u_\infty^*}(k, x) - x^*\| \leq \rho(k) \quad \text{for all } k \in \mathbb{N}.$$

- (ii) The finite horizon problems have the turnpike property at x^* in the following sense: for each $c \in (0, 1)$ there exists $\sigma_c \in \mathcal{L}$ such that for each optimal trajectory $x_{u_N^*}(k, x)$, $x \in X$ and all $N \in \mathbb{N}$ there is a set $\mathcal{Q}(x, N, c) \subseteq \{0, \dots, N\}$ with $\#\mathcal{Q}(x, N, c) \geq cN$ and

$$\|x_{u_N^*}(k, x) - x^*\| \leq \sigma_c(N) \quad \text{for all } k \in \mathcal{Q}(x, N, c).$$

- (iii) The optimal value functions V_N and V_∞ are continuous at x^* in the following uniform way: there is an open ball $\mathcal{B}_\varepsilon(x^*)$, $\varepsilon > 0$, around x^* and $\alpha \in \mathcal{K}_\infty$ such that for all $x \in \mathcal{B}_\varepsilon(x^*)$ and all $N \in \mathbb{N} \cup \{\infty\}$ the inequality

$$|V_N(x) - V_N(x^*)| \leq \alpha(\|x - x^*\|)$$

holds.

□

We remark that the fact that the bounds in (i) and (ii) are assumed to be independent of x may be restrictive if X is unbounded. However, since the interesting dynamics usually takes place in a bounded set X , we decided to make this assumption which considerably simplifies the subsequent arguments. A discussion of the turnpike property (ii) is provided in Section 3.3.

Lemma 3.2 (i) If Assumption 3.1 (i) and (iii) hold, then the equation

$$V_\infty(x) = J_K(x, u_\infty^*) + \beta^K V_\infty(x^*) + R_1(x, K) \quad (3.9)$$

holds with $|R_1(x, K)| \leq \alpha(\rho(K))$ for all sufficiently large $K \in \mathbb{N}$ and all $x \in X$.

(ii) For any $c \in (0, 1)$, if Assumption 3.1 (ii) and (iii) hold, then the equation

$$V_N(x) = J_K(x, u_N^*) + \beta^K V_{N-K}(x^*) + R_2(x, K, N) \quad (3.10)$$

holds with $|R_2(x, K, N)| \leq \alpha(\sigma_c(N))$ for all sufficiently large N , all $x \in X$ and all $K \in \mathcal{Q}(x, N, c)$.

Proof: (i) The dynamic programming principle yields

$$V_\infty(x) = J_K(x, u_\infty^*) + \beta^K V_\infty(x_{u_\infty^*}(K, x)).$$

Hence, (3.9) holds with $R_1(x, K) = \beta^K V_\infty(x_{u_\infty^*}(K, x)) - \beta^K V_\infty(x^*)$. Now choose $K \in \mathbb{N}$ so large that $\rho(K) < \varepsilon$ holds for ρ from Assumption 3.1 (i) and ε from Assumption 3.1 (iii). Then we obtain $|R_1(x, K)| \leq \beta^K \alpha(\|x_{u_\infty^*}(K, x) - x^*\|) \leq \alpha(\rho(K))$ and thus the assertion.

(ii) In the finite horizon case, the dynamic programming principle yields

$$V_N(x) = J_K(x, u_N^*) + \beta^K V_{N-K}(x_{u_N^*}(K, x)).$$

Hence, (3.10) holds with $R_2(x, K, N) = \beta^K V_{N-K}(x_{u_N^*}(K, x)) - \beta^K V_{N-K}(x^*)$. Now choose $N \in \mathbb{N}$ so large that $\sigma_c(N) < \varepsilon$ holds for σ_c from Assumption 3.1 (ii) and ε from Assumption 3.1 (iii). For any $K \in \mathcal{Q}(x, N, c)$ this implies $|R_2(x, K, N)| \leq \beta^K \alpha(\|x_{u_N^*}(K, x) - x^*\|) \leq \alpha(\sigma_c(N))$ and thus the assertion. \square

Lemma 3.3 If Assumption 3.1 (i)–(iii) holds, then for any $c \in (0, 1)$ the equation

$$J_K(x, u_\infty^*) = J_K(x, u_N^*) + R_3(x, K, N) \quad (3.11)$$

holds with $|R_3(x, K, N)| \leq \alpha(\rho(K)) + \alpha(\sigma_c(N))$ for all sufficiently large N , all $x \in X$ and all sufficiently large $K \in \mathcal{Q}(x, N, c)$.

Proof: The finite horizon dynamic programming principle implies that $u = u_N^*$ minimizes the expression $J_K(x, u) + \beta^K V_{N-K}(x_u(K, x))$. Together with the error term R_2 defined in the proof of Lemma 3.2 (ii) and $\tilde{R}_1(x, K, N) = \beta^K V_{N-K}(x_{u_\infty^*}(K, x)) - \beta^K V_{N-K}(x^*)$ this yields

$$\begin{aligned} J_K(x, u_N^*) + \beta^K V_{N-K}(x^*) &= J_K(x, u_N^*) + \beta^K V_{N-K}(x_{u_N^*}(K, x)) - R_2(x, K, N) \\ &\leq J_K(x, u_\infty^*) + \beta^K V_{N-K}(x_{u_\infty^*}(K, x)) - R_2(x, K, N) \\ &= J_K(x, u_\infty^*) + \beta^K V_{N-K}(x^*) - R_2(x, K, N) + \tilde{R}_1(x, K, N). \end{aligned}$$

Similar to the proof of Lemma 3.2 (i) one sees that $|\tilde{R}_1(x, K, N)| \leq \alpha(\rho(K))$ for all sufficiently large K .

Conversely, the infinite horizon dynamic programming principle implies that u_∞^* minimizes the expression $J_K(x, u_\infty^*) + \beta^K V_\infty(x_{u_\infty^*}(K, x))$. Using the error terms R_1 from the proof of Lemma 3.2 (i) and $\tilde{R}_2(x, K, N) = \beta^K V_\infty(x_{u_N^*}(K, x)) - \beta^K V_\infty(x^*)$ we obtain

$$\begin{aligned} J_K(x, u_\infty^*) + \beta^K V_\infty(x^*) &= J_K(x, u_\infty^*) + \beta^K V_\infty(x_{u_\infty^*}(K, x)) - R_1(x, K) \\ &\leq J_K(x, u_N^*) + \beta^K V_\infty(x_{u_N^*}(K, x)) - R_1(x, K) \\ &= J_K(x, u_N^*) + \beta^K V_\infty(x^*) - R_1(x, K) + \tilde{R}_2(x, K, N). \end{aligned}$$

As in the proof of Lemma 3.2 (ii) one obtains $|\tilde{R}_2(x, K, N)| \leq \alpha(\sigma_c(N))$. Together with the estimates for R_1 and R_2 from Lemma 3.2 this yields

$$\begin{aligned} |R_3(x, K, N)| &= |J_K(x, u_\infty^*) - J_K(x, u_N^*)| \\ &\leq \max\{|R_1(x, K)| + |\tilde{R}_2(x, K, N)|, |\tilde{R}_1(x, K, N)| + |R_2(x, K, N)|\} \\ &\leq \alpha(\rho(K)) + \alpha(\sigma_c(N)). \end{aligned}$$

\square

Theorem 3.4 If Assumption 3.1 (i)–(iii) holds and V_∞ is bounded on X , then the inequality

$$J_\infty^{cl}(x, \mu_N) \leq V_\infty(x) + \frac{\delta(N)}{1 - \beta}$$

holds for all sufficiently large $N \in \mathbb{N}$ with

$$\delta(N) = (2 + 2\beta)\alpha(\rho(\lceil N/2 \rceil - 2)) + 2\alpha(\sigma_{1/2}(N - 1)).$$

Proof: Picking $x \in X$ and abbreviating $x^+ := f(x, \mu_N(x))$, the dynamic programming principle and the definition of μ_N yield

$$g(x, \mu_N(x)) = V_N(x) - \beta V_{N-1}(x^+).$$

Moreover, again by the dynamic programming principle, the optimal control $u_{N-1}^{*,+}$ for x^+ coincides with $u_N^*(\cdot + 1)$. Thus, for any $K \in \mathcal{Q}(x, N, c)$ with $K \geq 1$ we obtain¹² $K - 1 \in \mathcal{Q}(x^+, N - 1, c)$ and $R_2(x, K, N) = R_2(x^+, K - 1, N - 1)$. Using (3.10) for N with x and K and with $N - 1$, x^+ and $K - 1$, respectively, yields

$$\begin{aligned} & V_N(x) - \beta V_{N-1}(x^+) \\ &= J_K(x, u_N^*) + \beta^K V_{N-K}(x^*) + R_2(x, K, N) \\ &\quad - \beta J_{K-1}(x^+, u_{N-1}^*) - \beta \beta^{K-1} V_{N-1-(K-1)}(x^*) - \beta R_2(x^+, K - 1, N - 1) \\ &= J_K(x, u_N^*) - \beta J_{K-1}(x^+, u_{N-1}^*) + (1 - \beta)R_2(x, K, N). \end{aligned}$$

A similar estimate using (3.9) yields

$$V_\infty(x) - \beta V_\infty(x^+) = J_K(x, u_\infty^*) - \beta J_{K-1}(x^+, u_\infty^{*,+}) + R_1(x, K) - \beta R_1(x^+, K - 1).$$

Note that here we cannot identify the two R_1 -terms because x^+ is computed using $u_N^*(0)$ and not $u_\infty^*(0)$, hence $u_\infty(\cdot + 1)$ does in general not coincide with the optimal control $u_\infty^{*,+}$ for x^+ .

Putting the three equations together and using Lemma 3.3 yields

$$g(x, \mu_N(x)) = V_\infty(x) - \beta V_\infty(x^+) + R_4(x, K, N).$$

with

$$\begin{aligned} R_4(x, K, N) &= (1 - \beta)R_2(x, K, N) - R_3(x, K, N) + \beta R_3(x^+, K - 1, N - 1) \\ &\quad - R_1(x, K) + \beta R_1(x^+, K - 1). \end{aligned}$$

From the bounds on the individual R_i -terms we obtain $|R_4(x, K, N)| \leq \delta(N)$, observing that for $c = 1/2$ the set $\mathcal{Q}(x, N, 1/2)$ contains at least one $K \geq \lceil N/2 \rceil - 1$ which we choose for obtaining the inequality for R_4 . From this we obtain

$$\begin{aligned} J_\infty^{cl}(x, \mu_N) &= \sum_{k=0}^{\infty} \beta^k g(x_{\mu_N}(k, x), \mu_N(x_{\mu_N}(k, x))) \\ &\leq \sum_{k=0}^{\infty} \beta^k \left(V_\infty(x_{\mu_N}(k, x)) - \beta V_\infty(x_{\mu_N}(k + 1, x)) + \delta(N) \right) \\ &= V_\infty(x) + \frac{\delta(N)}{1 - \beta} \end{aligned}$$

due to the boundedness of V_∞ . This proves the claim. \square

¹²Here we assume without loss of generality that the sets \mathcal{Q} are always chosen maximally.

3.3 The turnpike property

Theorem 3.4 heavily relies on the turnpike property defined in Assumption 3.1 (ii). By this property it is guaranteed that finite-horizon optimal trajectories are close to the optimal equilibrium x^* for a minimum number of time instants. The distance between the trajectory and x^* at these time instants reduces as N increases. The typical behaviour of an optimal trajectory is as follows: The trajectory approaches the optimal equilibrium (the turnpike), stays there for some time and turns away in the end. For infinite-horizon optimal trajectories (not considered here) the turnpike property ensures convergence to the optimal equilibrium.

Turnpike properties do not only play an important role in our analysis (see also Grüne (2013) and Grüne and Stieler (2014)), but are indeed a well known and investigated research area in mathematical and economic literature since the works of von Neumann (1945) and Dorfman et al. (1987). The latter originally appeared in 1958 and was eponymous for the phenomenon. Turnpike properties have been studied for different types of models and in different settings, such as discrete-time and continuous-time, discounted and undiscounted, finite- and infinite-horizon, stochastic and deterministic models.

In economic literature one of the main reasons for establishing turnpike properties and searching for sufficient conditions for them to hold is to provide insights to specific models and to guarantee a certain behaviour of the optimal trajectories/solutions. Besides those mentioned above, one of the dominant figures in economic turnpike theory was Lionel W. McKenzie, see, e.g., McKenzie (1986).

In mathematical literature, turnpike phenomena are investigated rather for different types of optimal control problems than for specific models, see e.g., Zaslavski (2006) and Zaslavski (2014) for a collection of turnpike theorems or Damm et al. (2014) for an exponential turnpike theorem (i.e. σ_c from Assumption 3.1 (ii) is exponentially decaying in N). As in economic literature one is interested in the behaviour of optimal solutions. Beyond this, turnpike properties are a key ingredient to establish performance estimates and stability of MPC closed-loop solutions, see Grüne (2013), Grüne and Stieler (2014) and Theorem 3.4.

There are many references on sufficient conditions for turnpike properties in case of discounted optimal control problems. For example in Scheinkman (1976) the turnpike property is proved assuming that the discount factor β is sufficiently close to 1. In Brock and Scheinkman (1976) an explicit bound for the discount factor or, more precisely, the discount rate in the considered continuous time model, is given. In (Zaslavski; 2014, Section 2.25) we can find a similar assumption in Theorem 2.49 ff.: It is required that for some arbitrary discounting sequence the discounting factors do not change too fast. In our setting this is equivalent to assuming that β is close to 1. The fact that turnpike results in the literature are primarily for discount factors β close to 1 does not affect their usefulness for our purpose. On the contrary, this is precisely the situation in which we need them for our convergence analysis.

Beyond our theoretical analysis, numerical results indicate that the performance of NMPC may indeed deteriorate if the turnpike phenomenon does not occur for sufficiently small N . This is due to the fact that the numerical effort for solving the finite horizon optimal control problem (2.8) grows with N and for too large horizons N optimization algorithms may fail to deliver reliable solutions due to numerical problems. Thus, if the optimal trajectories

only exhibit the turnpike property for N for which the problem is numerically infeasible, NMPC may not produce near optimal solutions. This is, for instance, what happens in the example in Section 4.2. Fortunately, there is a remedy for this problem, by using the infinite horizon optimal steady state x^* as terminal constraint when solving (2.3), as proposed in Angeli et al. (2009); Angeli and Rawlings (2010); Diehl et al. (2011). More precisely, in each iteration we perform the maximization in (2.3) only over those control sequences $u^i(\cdot)$ for which $x(N) = x^*$ holds. Proceeding this way forces the trajectories to exhibit the turnpike phenomenon for smaller horizons N and can thus improve performance. We note that in this paper such terminal constraints were only used in the example in Section 4.2; all other computations in this paper were performed without using terminal (steady state) constraints when solving (2.3). Further economic examples in which terminal constraints turn out to be beneficial can be found in Bréchet et al. (2012).

Finally, it would be very interesting to extend the analysis from the previous section to problems which exhibit a more complicated limit dynamics than optimal equilibria. The effects which arise in this case are, however, not yet fully understood and will be subject of future research.

4 One dimensional examples

In this section we describe the application of our NMPC algorithm to a selection of one dimensional optimal control problems. For some one dimensional models the use of sophisticated numerical algorithms is not really necessary, because these problems can usually be solved with high precision in a reasonable amount of time with numerous procedures. Nevertheless, such (numerically) simple problems are important for verifying the accuracy of numerical procedures. Our first example will illustrate this, because for this problem the exact solution is known, hence the accuracy of the respective methods can be compared directly. For the subsequent examples the exact solution is not known, hence numerical methods are necessary for their analysis.

4.1 The basic growth model

We start our numerical investigations with a basic growth model in discrete time, which goes back to Brock and Mirman (1972) and has triggered extensive research in the RBC (real business cycle) literature.¹³ This model has also been used as a test example for many numerical algorithms, see, e.g., Santos and Vigo-Aguiar (1995, 1998, Sect. 4) and Grüne and Semmler (2004). The problem is a discrete time maximization problem of type (2.3)–(2.4) with the payoff function and dynamics given by

$$g(x, u) = \ln(Ax^\alpha - u) \quad \text{and} \quad x(t+1) = u(t).$$

Hereby, Ax^α is a production function with constants $A > 0$, $0 < \alpha < 1$, capital stock x and control variable u . The difference between output and next period's capital stock (given by u) is consumption. The exact solution to this problem is known (see Santos and

¹³For a stochastic version of this basic growth model, see Section 7.

Vigo-Aguiar (1998)) and is given by

$$V(x) = B + C \ln x, \quad \text{with } C = \frac{\alpha}{1 - \alpha\beta} \text{ and } B = \frac{\ln((1 - \alpha\beta)A) + \frac{\beta\alpha}{1 - \beta\alpha} \ln(\alpha\beta A)}{1 - \beta}.$$

The unique optimal equilibrium for this example is given by $x^* = 1/\sqrt[\alpha]{\beta\alpha A}$ and as we specify parameters $A = 5$, $\alpha = 0.34$ and $\beta = 0.95$ for our numerical tests, we have $x^* \approx 2.067$.




Figure 4.1: Closed loop trajectory (solid) and open loop trajectories (dashed) for the growth model for $N = 5$ and $x_0 = 5$ (left) and $V(x_0) - J_\infty(x_0, \mu_N)$ for $x_0 = 5$, $N = 2, \dots, 15$ (right), with μ_N the NMPC controller for time horizon N .

We expect that the NMPC algorithm computes closed loop solutions such that the system is steered into a neighborhood of the optimal equilibrium, and this is indeed the observed behavior in Figure 4.1 (left; the closed loop solution is depicted as solid line with circles).¹⁴ Moreover, we see that the open loop trajectories (shown as dashed lines) exhibit the turnpike property (cf. Section 3.3 and also Figure 7.11).

The second interesting aspect to explore is whether the NMPC closed loop trajectory maximizes the given objective function. In Figure 4.1 (right) we compare the exact optimal value $V(x_0)$ to the return generated by the NMPC algorithm with different N for initial value $x(0) = x_0 = 5$, i.e. we compute

$$V(x_0) - J_\infty(x_0, \mu_N), \quad J_\infty(x_0, \mu_N) := \sum_{t=0}^{\infty} \beta^t g(x(t), \mu_N(x(t))),$$

where μ_N denotes the NMPC controller for horizon N . Figure 4.1 shows, that the difference $V(x_0) - J_\infty(x_0, \mu_N)$ converges to zero exponentially fast for $N \rightarrow \infty$ and hence, for increasing N , the NMPC generated return $J_\infty(x_0, \mu_N)$ approximates the optimal value $V(x_0)$ arbitrarily well.

¹⁴The MATLAB code for this example can be found on the webpage http://num.math.uni-bayreuth.de/de/publications/2015/gruene_semmler_stieler_using_NMPC_2015

4.2 The basic DSGE model

Next we describe an extension of the basic growth model of section 3.1, where we now introduce also a labor choice in the preferences, as is commonly done in DSGE (dynamic stochastic general equilibrium) models. As to the formulation of preferences we follow Aruoba et al. (2006) and Parra-Alvarez (2012) but we do not pursue the strategy of linearization about the steady state.¹⁵ We consider the non-stochastic variant of it which has one state variable and two control variables, and attempt to find global solutions.

The model, in continuous time form, looks as follows

$$g(x, u) = \frac{(u_1(1 - u_2)^\psi)^{(1-\gamma)}}{1 - \gamma}$$

with dynamics

$$\frac{d}{dt}x(t) = (r(t) - \rho)x(t) + w(t)u_2(t) - u_1(t) .$$

with u_1 consumption, u_2 labor effort, x as capital stock, $r(t) = \alpha Ax(t)^{\alpha-1}u_2(t)^{1-\alpha}$ the return on capital, $w(t) = (1 - \alpha)Ax(t)^\alpha u_2(t)^{-\alpha}$ the wage rate, derived from a production function such as $Ax(t)^\alpha u_2(t)^{1-\alpha}$.

The following are standard parameters for this kind of model, see Aruoba et al. (2006) and Parra-Alvarez (2012). We set $A = 1$, $\alpha = 0.4$, $\rho = 0.0196$, $\gamma = 3.85$ and $\psi = 1.8011$. The discount rate for the dynamic decision problem is taken as $\delta = 0.010584$. The steady state values are¹⁶ $x^* \approx 23.03$, $u_1^* \approx 1.74$, $u_2^* \approx 0.31$.

In this example it turns out that for reasonable values of N the NMPC algorithm is not able to find the optimal steady state x^* (cf. the discussion in Section 3.3). As a remedy, we use $x^* = 23.03$ as a terminal constraint for the open loop optimization, i.e., we solve (2.8) under the additional constraint $x(N) = x^*$. In Figure 4.2 the resulting closed loop solution and open loop solutions are shown. Choosing x^* as a terminal constraint can lead to wrong or infeasible solutions in case x^* is not asymptotically stable. Since our model and parameters are defined as in Parra-Alvarez (2012, Proposition 3.2), optimal consumption and optimal labor effort in feedback form are known. The resulting differential equation has two equilibria, $x = 0$ and $x = x^*$. Our numerical experiments show that the dynamics converge to x^* for $x(0) > 0$, indicating asymptotic stability of x^* . Hence, using this equilibrium as a terminal constraint and thus forcing the solution to converge to x^* is justified.

We remark that this example is the only one in this paper where terminal constraints were imposed. For all the other examples NMPC without terminal constraints turned out to be

¹⁵The numerical solution techniques as they are implemented by DYNARE use mostly local techniques where an approximation is taken around the deterministic steady state. DYNARE can also solve dynamic decision models globally by using the deterministic steady state as terminal condition. Algorithms based on the perturbation method have been developed. These algorithms build on a Taylor series expansion of the agents' policy functions around the steady state of the economy and a perturbation parameter. In earlier literature one has used the first term of this series. Since the policy functions resulting from a first order approximation are linear and many dynamic models display behavior that is close to a linear law of motion, the approach became quite popular under the name of linearization. Judd and Guu (1997) extended the method to compute the higher-order terms of the expansion, see also Collard and Juillard (2001).

¹⁶For details see Parra-Alvarez (2012)




Figure 4.2: Closed loop solution only (left) and closed loop and open loop solutions (right), in both cases steady state of capital stock $x^* = 23.03$ used as terminal condition.

capable of finding sufficiently accurate solutions. We also mention that using the steady state x^* as terminal constraint does not mean that we perform linearization around x^* . The optimization is still performed for the full nonlinear model, the terminal constraint only serves as a “guideline” for the optimizer in order to find a solution approaching the optimal steady state.

5 Two dimensional examples

Subsequently, we describe two two-dimensional problems which we have solved using our algorithm. We first turn to a model with one control variable and two domains of attraction, where the study of the dynamics have been quite a challenge for research in economics, since here one expects the separation of domains of attraction given not by threshold points but threshold lines (Skiba lines). Furthermore, a two-dimensional model is added with two control variables where it is of economic relevance to track the finite time path of the solution trajectories.

5.1 A 2d model with multiple domains of attraction

The following problem from Haunschmied et al. (2003) is a 2d variant of an investment problem of the firm where the control variable is the change of investment rather than investment itself as in the usual case. The payoff function is here given by

$$g(x, u) = R(x_1) - c(x_2) - v(u)$$

and dynamics by

$$\frac{d}{dt} \begin{pmatrix} x_1(t) \\ x_2(t) \end{pmatrix} = \begin{pmatrix} x_2(t) - \sigma x_1(t) \\ u(t) \end{pmatrix}.$$

In the above model $R(x_1) = k_1\sqrt{x_1} - x_1/(1 + k_2x_1^4)$ is a revenue function of the firm with a convex segment due to increasing returns. $c(x_2) = c_1x_2 + c_2x_2^2/2$ denotes adjustment costs of investment and $v(u) = \alpha u^2/2$ represents adjustment costs of the change of investment.

The convex segment in the payoff function just mentioned is likely to generate two domains of attraction.

The following parameters will be used for our calculations: $\sigma = 0.25$, $k_1 = 2$, $k_2 = 0.0117$, $c_1 = 0.75$, $c_2 = 2.5$, $\alpha = 12$ and discount rate $\delta = 0.04$. Generating the vector field we have obtained the results shown in the following Figure 5.3.¹⁷




Figure 5.3: Vector field of the model with multiple attractors

The vector field shows clearly two domains of attraction,¹⁸ one at roughly $x_1^* = 0.5$, $x_2^* = 0.2$ and the other roughly at $x_1^* = 4.2$, $x_2^* = 1.1$. The vector field shows that there is a bifurcation of the dynamics, where the trajectories go either to the low level steady state or high level steady states. This bifurcation line has been called a Skiba curve in Haunschmied et al. (2003). The above vector field was generated by DP¹⁹ and serves as a benchmark for assessing the quality of the NMPC solutions.

Next, we want to test if we can replicate the two domains of attraction for a finite decision horizon by using NMPC. Therefore, we choose different initial values from both domains of attraction and run the NMPC algorithm for different horizons N . Figure 5.4 shows the resulting phase plots for selected N . We observe that for small N , e.g. $N = 10$, all NMPC trajectories converge to the left equilibrium. If we increase N , e.g. $N = 35$, the trajectories of some of the initial values converge to the second equilibrium. For $N = 50$, as in Figure 5.3, we observe the existence of a Skiba curve, i.e. trajectories resulting from initial values right (left) of the curve converge to the equilibrium on the right (left).

Note that if agents have different decision horizons, this might actually give rise to different long run steady states.²⁰ A further study of the numerical challenges that the above example poses are explored in more general terms in Section 7.

¹⁷Figure 5.3 was generated through DP, with time step $h = 1/20$ and 101 control values \tilde{U} covering $U = [-1, 1]$. For the details of solving the model through DP and generating the vector field, see Grüne and Semmler (2004).

¹⁸A DSGE type model with two domains of attraction and regime changes can be found in Farmer et al. (2009). In future work of the authors the latter type of model is intended to be explored.

¹⁹See Grüne and Semmler (2004)

²⁰Thus, if one interprets the finite decision horizon problem as a form of the inattentiveness problem in the sense of Sims, a too short a decision horizon may be misleading.




Figure 5.4: Phase plots of NMPC trajectories of the model of multiple domains of attraction for $N = 10$ (upper left), $N = 35$ (upper right) and $N = 50$ (lower)

5.2 A 2d growth model with non-renewable resources

As the next example of this section we study a tracking problem. We consider a growth model with the extraction of a non-renewable resource as discussed in Greiner and Semmler (2008, Ch. 14)²¹. The model is as follows:

$$g(x, u) = U(u_1)$$

$$\frac{d}{dt} \begin{pmatrix} x_1(t) \\ x_2(t) \end{pmatrix} = \begin{pmatrix} F(x_1(t), u_2(t)) - u_1(t) - \sigma x_1(t) \\ -u_2(t) \end{pmatrix}$$

and control constraints

$$\begin{aligned} 0.01 &\leq u_1(t) \leq 3.0 \\ 0.01 &\leq u_2(t) \leq 2.0. \end{aligned}$$

The model captures the extraction of the non-renewable resource needed for production. It posits that there is utility from consumption, $U(u_1) = (u_1(t)^{1-\sigma}) / (1-\sigma)$. The production

²¹This is a sub-model of our high dimensional model in Section 6




Figure 5.5: Growth model with non-renewable resource, closed loop solid, open loop dashed, initial value $x(0) = (1, 10)^T$

function includes extracted exhaustible resources $F(x_1(t), u_2(t)) = x_1(t)^\beta u_2(t)^{1-\beta}$. Here x_1 and x_2 are the capital stock and the stock of resources, where the stock of extracted resources is constrained by $x_2(t) \leq x_2(0)$ (due to the positive lower bound on u_2) with $x_2(0)$ the initial stock of the exhaustible resource. The control variables u_1 and u_2 are consumption and the extracted resource per unit of time.

We set the parameters to $\rho = 0.03$, $\sigma = 0.5$, $\beta = 0.7$, and $\delta = 0.05$. Figure 5.5 shows how the dynamics of the optimally growth model with the non renewable resource behave.

Initially the non-renewable resource is assumed to be $x_2(0) = 10$, and $x_1(0) = 1$. As one can observe the path for the non-renewable resource declines to zero, using an optimal extraction strategy, and the capital stock first builds up and then, since the non-renewable resources becomes exhausted, also declines to zero. The results reported here are consistent with the literature starting with Hotelling.²² The results in the literature typically replicate Hotelling's presumption, namely that the resource is optimally extracted until there is no more resource in situ. Our NMPC algorithm lets us conveniently track the paths of the resource and capital stock in a model with decisions on a finite time horizon. This will prove to be a useful procedure to study the case when higher dimensional models are explored. This is done next.

6 A five dimensional problem

The following example resembles the integrate assessment model of growth and climate change²³ that includes the non-renewable energy resource of Section 5.2, an alternative

²²Yet, in Hotelling there is no capital stock and consumption included in the resource extraction model.

²³See Nordhaus (2009).

energy resource, such as renewable energy, and an adaptation policy. Here also the issue of a tracking problem may come up. In the preferences of households it considers, beside consumption, damages arising from the use of non-renewable energy and an investment effect reducing those damages.²⁴ Here too, we model the dynamic decision problem over a finite horizon, with three control variables and five state variables. It is formulated as continuous time problem such as

$$\max_{c(t), e_p(t), u(t)} \int_0^{\infty} e^{-(\rho-n)t} \left(\frac{\left(c(t) \cdot (\alpha_2 e_p(t))^\eta \left(M(t) - \widetilde{M} \right)^{-\varepsilon} (\nu_2 g(t))^\omega \right)^{1-\sigma} - 1}{1-\sigma} \right) dt$$

s.t.

$$\dot{k}(t) = A(A_n k(t) + A_r u(t))^\alpha (\nu_1 g(t))^\beta - c(t) - e_p(t) - (\delta_k + n)k(t) - u(t) \cdot (\psi R(t))^{-\tau}$$

$$\dot{g}(t) = i_p^f + \alpha_1 e_p(t) - (\delta_g + n)g(t) + i_p^b$$

$$\dot{b}(t) = (r - n)b(t) - (1 - \alpha_1 - \alpha_2 - \alpha_3)e_p(t) + i_p^b$$

$$\dot{R}(t) = -u(t)$$

$$\dot{M}(t) = \gamma u(t) - \mu \left(M(t) - \kappa \widetilde{M} \right) - \theta (\nu_3 g(t))$$

The state variables k , g , b , R , M denote the stock of private capital, the stock of public capital, the outstanding debt, the stock of non-renewable resource and the concentration of greenhouse gases (GHGs). The control variables c , e_p , u represent consumption, amount of resources absorbed by the public sector and the resource extraction rate. The model of Section 5.2 is present here in the dynamics for \dot{R} and \dot{k} , though in the production function in \dot{k} , in the first term, output is produced by non-polluting energy $A_n k$ and polluting energy $A_r u$, using the extracted resource u . Moreover, a fraction of public capital g is used for the production of output. The last term in the capital stock equation represents the extraction cost of the resource, which is inverse to the available resource R .

The second equation, \dot{g} , represents the evolution of public capital, funded by a fraction of the tax revenue available for investment on infrastructure. The constant terms are parameters.

The evolution of public debt is denoted by the \dot{b} equation with the first term the interest payment, and the second term $(1 - \alpha_1 - \alpha_2 - \alpha_3)e_p$ the excess of tax revenue over spending, the fiscal surplus, to be used for debt redemption. The last equation represents in the first term the pollution effect from using non-renewable energy, the second term is the natural dissipation of GHGs M and the last term is an expression for adaption efforts, also using up a fraction of public capital.

A list of parameters for the model is provided in Appendix A. Since we assume here a growing population, with growth rate n , with all variables written in per capita form, the discount rate is $\rho - n$. This example is a nonlinear dynamic decision model of higher dimension. We hereby will also study the impact of state and control constraints on the solution of the model.

In all the Figures we used a sampling rate of $T = 0.1$ and posited constraints such k , g , b ,

²⁴The different components of the subsequent model can be found in Greiner and Semmler (2008), Greiner et al. (2007) and Greiner et al. (2014).

$R, M, c, e_p, u \geq 0$. First, we simulated the NMPC solutions for state constraints $k, g, b, R, M \leq 12$ and control constraints $c, e_p \leq 2, u \leq 4$. It turns out that the state constraints are not active except for the state b . For this component, however, we observe that the closed loop solution violates the constraint $b \leq 12$, see graph in Figure 6.6.




Figure 6.6: States (left) and controls (right) of the five dimensional model with constraint $b \leq 12$ and $N = 15$.

At the time when b hits the upper bound the control e_p is increased trying to keep b feasible but it does not succeed. This causes a kink in the graph of k , see the right figure.

A natural question, arising from the observed behaviour, is whether the states k, g, R, M converge if the upper bound of b is omitted. Then, b exhibits unbounded growth and – as can be seen in Figure 6.7 – all other state variables converge. Moreover, no numerical




Figure 6.7: State trajectories (left) of the five dimensional model with unrestricted state b and $N = 15$ and related controls (right) on the time interval $t \in [0, 20]$.

difficulties, indicated by kinks or zig-zags, are observed. Note that in our dynamics for b , there is no feedback from the level of debt to the fiscal surplus. For a detailed discussion of this issue see Greiner et al. (2007), and the discussion below. We can interpret Figure 6.7 as follows: If the debt b is unbounded, this yields a strategy of immediately exploiting non-

renewables (u increases $\Rightarrow R$ is decreasing quickly) which leads to an increased emission of GHGs (M). Then, consumption c is high at the beginning but then converges to quite a low level. Public and private capital stagnate at the level of the respective initial stocks.

In any reasonable model, however, one would not like to have the debt variable b to be unbounded. Debt sustainability is a major issue in economic modeling. Hence, we intend to explore why the state constraint $b \leq 12$ is not maintained and how this can be fixed. A possible explanation is that the optimization horizon was chosen too small such that the optimizer does not “see” the bound at times when it is still possible to find controls that remain in the state feasible for future times. This fact can be formalized via stationarity of feasible sets, cf. Chapter 5 of Kerrigan (2000) or equivalently via exit sets, cf. Section 8.2 of Grüne and Pannek (2011). Accordingly, the problem should vanish if N is increased and/or the bound on b is reduced. Figure 6.8 shows that this is indeed what happens if we increase N to 30 and additionally reduce the bound on b .




Figure 6.8: State trajectories (left) of the five dimensional model with $b \leq 1.0$, $N = 30$ and related controls (right) on the time interval $t \in [0, 100]$.

Numerical experiments revealed that the size of the horizon N needed to ensure constraint satisfaction grows with the upper bound on b , i.e., N has to be increased when the upper bound on b is raised. Figure 6.9 shows the simulation of the model with $b \leq 2.0$ and $N = 50$.

In this context, it is interesting to compare the optimal strategies without and with bound on the debt in Figure 6.9: as b hits the bound, the public sector is forced to use more resources (e_p is raised). In this case, with an increase of revenue e_p , the fiscal surplus in the dynamics for b (the debt equation) becomes endogenous, which allows the debt to be bounded. This also permits the public capital g to increase. The consumption c is reduced until $t \approx 80$ in order to avoid reduction of the private capital k and in favor of a higher constant level of consumption in a future time period. In contrast to the Figures 6.7 and 6.8, the strategy chosen in Figure 6.9 ensures higher levels of capital (private and public) and consumption in the long run if one is willing to accept cuts at the beginning.




Figure 6.9: State trajectories (left) of the five dimensional model with $b \leq 2.0$, $N = 50$ and related controls (right).

7 Some Implementational Aspects

In the previous sections various examples of finite horizon optimal control problems and their numerical results have been presented. All simulations were either computed by a MATLAB routine²⁵ (for less demanding examples) or by a C++ software package²⁶ (for more complicated examples).

The main part of implementing the NMPC algorithm consists of solving the optimization problem in step (2) of the NMPC algorithm. This is accomplished by transforming the given optimization problem into standard form

$$\begin{aligned} \min_{z \in \mathbb{R}^{n_z}} f(z) \\ \text{s.t. } G(z) = 0 \text{ and } H(z) \leq 0. \end{aligned} \quad (7.12)$$

To this end, we need to decide which variables should be chosen as optimization variables z . In Grüne and Pannek (2011, Chapter 10), three different approaches to that problem (also referred to as *discretization*) are proposed²⁷:

- 1) In *full discretization* not only the control values $u_{k,i}$, $k = 0, \dots, N - 1$, but also the states $x_{k,i}$, $k = 0, \dots, N$, are considered as optimization variables. The inclusion of the states requires additional equality constraints which ensure that the trajectory

²⁵available at www.nmpc-book.com

²⁶see www.nonlinearmpc.com

²⁷For the sake of comprehensibility we only consider systems in discrete time. In case of continuous systems, we replace the system dynamics by a numerical approximation, cf. (2.6).

satisfies the system dynamics. This leads to the following definitions in (7.12):

$$z := (u_{0,i}^T, \dots, u_{N-1,i}^T, x_{0,i}^T, \dots, x_{N,i}^T)^T, \quad f(z) := - \sum_{k=0}^{N-1} \beta^k g(x_{k,i}, u_{k,i}),$$

$$G(z) := \begin{pmatrix} *1 \\ x_{0,i} - x_0 \\ x_{1,i} - \varphi(x_{0,i}, u_{0,i}) \\ \vdots \\ x_{N,i} - \varphi(x_{N-1,i}, u_{N-1,i}) \end{pmatrix}, \quad H(z) := (*2),$$

where $*_1$ and $*_2$ denote possible pre-existing constraints.

- 2) *Recursive discretization* describes the approach to decouple the system dynamics from the optimization problem, i.e. only the control values $u_{k,i}$ are optimization variables whereas the system dynamics are computed outside the optimization. Since the optimizer requires information about the system and vice versa, both components need to communicate: The optimizer sends the initial value and the control values to the system dynamics which in turn sends the corresponding states that are needed in order to evaluate the objective function $f(z)$. Figure 7.10 shows the exchange of information schematically.




Figure 7.10: Exchange of information between the optimization problem and the system dynamics

Consequently, in (7.12) after recursive discretization we have

$$z := (u_{0,i}^T, \dots, u_{N-1,i}^T)^T, \quad f(z) := - \sum_{k=0}^{N-1} \beta^k g(x_{k,i}, u_{k,i}),$$

$$G(z) := (*1), \quad H(z) := (*2),$$

with $*_1$ and $*_2$ as above.

- 3) The third technique — called *shooting discretization* — includes some of the states $x_{k,i}$ as optimization variables, but in contrast to full discretization just for some of the $k \in \{0, \dots, N-1\}$ and possibly not for all components. Certainly, all $u_{k,i}$ are chosen as optimization variables as well. As in full discretization we need to impose additional equality constraints for the states which are optimization variables.

Obviously, a main disadvantage of technique 1) is the high dimensionality of the resulting optimization problem. At the same time, the optimizer is given the full information about the dynamics which is an advantage for iterative solvers, e.g., for obtaining good initial guesses. Moreover, the special structure of the resulting fully discretized optimization problem can be used in order to simplify the problem to be solved (using a technique called *condensing*, see Grüne and Pannek (2011, Sec. 10.4)).

In recursive discretization, the optimization problem has minimal dimension but information about the trajectories can hardly be used within the optimization. In addition, the external computation of the trajectories may lead to numerical instability due to the sensitive dependence of the values $x_{k,i}$ on the control values $u_{k,i}$: Even a small deviation of one of the $u_{k,i}$ may result in a large deviation of the trajectory and hence of the objective function.

Shooting discretization can be seen as an attempt to reduce the dimension of the fully discretized problem as much as possible without losing stability and useful information about the trajectories. For a detailed discussion on the three techniques see Grüne and Pannek (2011, Sec. 10.1). Regarding the software we used for the examples in this paper, the MATLAB routine we have used is based upon recursive discretization and uses the `fmincon` optimization routine for solving (7.12) while the C++ software is able to perform each of the discretization techniques and has links to various optimization packages which can be selected for solving (7.12).

Now that we have defined the optimization problem, we might expect difficulties whenever (7.12) is nonlinear and nonconvex. In this case, the optimization algorithm may only find a local optimum which does not need to be a global optimum or the optimizer may not be able to find an optimum, at all. While such difficulties did not occur in the examples in this paper, when interpreting the outcome of an NMPC algorithm one should always be aware that this may happen. Often, one can avoid such situations e.g. by adding constraints or choosing the initial guess of the optimizer carefully.

The optimization horizon N plays an important role in NMPC. As pointed out in Section 2, an approximation of the infinite horizon optimal trajectories can only be expected if N is sufficiently large. In Section 5.1 we already illustrated the effect of varying N in the presence of multiple optimal equilibria. Recall that in this example we needed to increase N to about 50 in order to obtain the correct domains of attraction. In order to explain why this happens and also in order to illustrate the turnpike property as the mechanism for the approximation property of NMPC (cf. the discussions in Sections 2.2 and 4.1), we reconsider the example from Section 4.1. In Figure 7.11 we show the optimal open loop trajectories starting in $x_0 = (3, 0.75)^T$ for different N .

This figure shows that the open loop trajectories are attracted by the lower left equilibrium for small N and by the upper right equilibrium for larger N , i.e., for too small N the optimizer does not “see” the proper optimal equilibrium. Moreover, we can observe the turnpike property for $N \geq 50$: the larger N , the closer the trajectories approach the optimal equilibrium (indicated by the “+” in the upper right corner of the figure) and the longer they stay in its neighborhood.

While it seems that increasing N is often a good strategy in order to obtain a good approximation of the infinite horizon optimal solutions, we want to point out that large horizons




Figure 7.11: Open loop trajectories for the model from Section 5.1 for $x_0 = (3, 0.75)^T$ and various $N \leq 45$ (dashed) and $N \geq 50$ (solid).

increase the dimensionality in (7.12) on the one hand and, on the other hand, may cause numerical problems (as shown in Grüne and Pannek (2011) for the inverted pendulum). Hence, there is a tradeoff between good approximation, numerical effort and numerical accuracy which implies that a judicious choice of N can only be found if all these effects are taken into account.

Summarizing, being aware of possible pitfalls, users of NMPC software should always interpret the obtained results with some care and consider a series of numerical experiments with different parameters in order to verify the validity of their results.

8 Outlook on NMPC for stochastic problems

In this section, by returning to the low dimensional problem of Section 3.1, we present some experimental results for NMPC applied to stochastic problems using the certainty equivalence principle. As we are not aware of any theoretical performance analysis for this approach in the context of NMPC modeling, we decided to present these promising results as an outlook and motivation for future research. Due to the fact that the control generated by the NMPC algorithm is in feedback form, the basic concept is easily extended to stochastic problems of the type

$$V(x_0) = E \left(\max_{u \in \mathcal{U}} \sum_{k=0}^{\infty} \beta^k g(x(k), u(k)) \right) \quad (8.13)$$

with the discrete time stochastic dynamics

$$x(k+1) = \varphi(x(k), u(k), z_k), \quad x(0) = x_0 \in \mathbb{R}^n, \quad (8.14)$$

where the z_k are i.i.d. random variables. Again, this problem could a priori be given in discrete time or it could be given as the time discretization of a continuous time stochastic optimal control problem with dynamics governed by an Itô-stochastic differential equation, see Camilli and Falcone (1995).

From a computational point of view, the main difficulty in stochastic NMPC is the efficient solution of the corresponding finite horizon problem (2.8) which now becomes a stochastic optimal control problem whose solution is computationally considerably more expensive than in the deterministic case. While some MPC approaches in the literature indeed solve stochastic optimal control problems (see, e.g., Couchman et al. (2006) or Cannon et al. (2009) and the references therein), in this paper we follow the simpler certainty equivalence approach similar to (Bertsekas; 2005, Section 6.1) which does in general not compute the true stochastic optimum but in the case of stochastic perturbations with low intensities may still yield reasonably good approximately optimal results. To this end, we replace the stochastic dynamics by its expected counterpart

$$x^e(k+1) = E\left(\varphi(x^e(k), u(k), z_k)\right), \quad x^e(0) = x_0 \in \mathbb{R}^n \quad (8.15)$$

and in each iteration instead of (2.8) we solve

$$\max_{u \in \mathcal{U}} \sum_{k=0}^{N-1} \beta^k g(x^e(k), u(k)). \quad (8.16)$$

Note that we only use (8.15) in order to solve (8.16) in Step (2) of the NMPC algorithm. In Step (3) we simulate the closed loop using the original stochastic dynamics (8.14) with z_k realized by appropriate random numbers.

We illustrate the performance of this approach by a two dimensional stochastic version of the Brock and Mirman (1972) model of Example 4.1. Here the 1d model from Example 4.1 is extended using a second variable modelling a stochastic shock. The model is given by the discrete time equations

$$\begin{aligned} x_1(k+1) &= x_2(k)Ax_1(k)^\alpha - u(k) \\ x_2(k+1) &= \exp(\rho \ln x_2(k) + z_k) \end{aligned}$$

where A, α and ρ are real constants and the z_k are i.i.d. random variables with zero mean. The payoff function in (2.3) is again $g(x, u) = \ln u$.

In our numerical computations we used the parameter values $A = 5$, $\alpha = 0.34$, $\rho = 0.9$ and $\beta = 0.95$ and z_k are i.i.d. Gaussian random variables with zero mean and variance $\sigma^2 = 0.008^2$. Using that $E(\exp(a + z_k)) = \exp(a + \sigma^2/2)$, the model used for the open loop optimization is given by

$$\begin{aligned} x_1^e(k+1) &= E(x_2^e(k)Ax_1^e(k)^\alpha - u(k)) = x_2^e(k)Ax_1^e(k)^\alpha - u(k) \\ x_2^e(k+1) &= E(\exp(\rho \ln x_2^e(k) + z_k)) = \exp(\rho \ln x_2^e(k) + \sigma^2/2). \end{aligned}$$

Following the computations in Santos and Vigo-Aguiar (1995) the optimally controlled dynamics is given by $x_1(k+1) = \alpha\beta Ax_2(k)x_1(k)^\alpha$. From this equation and the equation for $x_2(k+1)$, above, one can derive equations for the steady state values of $E(\ln x_1)$ and $E(\ln x_2)$ which transformed to the original exponential variables yield the expected equilibria $x_2^{e,*} = \exp(\sigma^2/(2(1-\rho^2))) \approx 1.000168$ and $x_1^{e,*} = (\alpha\beta A)^{\frac{1}{1-\alpha}} (x_2^{e,*})^{\frac{1}{1-\alpha^2}} \approx 2.067739$.

Figure 8.12 (left) shows the two components of a typical closed loop trajectory (solid) starting in $x_0 = (2, 1)^T$, along with the optimal open loop trajectories in each iteration (dashed). In order to measure the quality of the closed loop solutions, we have measured the average distance of the first component of the closed loop trajectory from the expected equilibrium. For each of these measurements an approximation $\tilde{E}(x_1(k))$ of this average was computed by a Monte-Carlo simulation using two trajectories starting in the optimal equilibrium $x^{e,*}$ with length 1000 and antithetic random numbers. Figure 8.12 (right) shows that the results improve with growing optimization horizon N until about $N = 8$, after which the errors caused by the Monte-Carlo simulation and the certainty equivalence approach become visible. Despite these errors, the simulations demonstrate that the NMPC approach is very well suited to compute approximately optimal trajectories also for stochastic problems with a reasonably small error.




Figure 8.12: Closed loop trajectory (solid) and open loop trajectories (dashed) for the growth model for $N = 5$ and $x_0 = (2, 1)^T$ (left) and equilibrium deviation $|\tilde{E}(x_1(t)) - x_1^{e,*}|$ for $x_0 = x^{e,*}$, $N = 3, \dots, 15$ (right), with μ_N the NMPC controller for time horizon N .

9 Conclusion

In the economic literature, there have been complicated dynamic models in various economic areas with optimizing agents, in part becoming a challenge to commonly used numerical techniques. In this paper we have demonstrated that NMPC provides an efficient way of numerically solving dynamic decision problems for some of the proto-type models of some of those areas to find global solutions. Our numerical findings are supported by a theoretical convergence analysis relying on the turnpike property of optimal solutions. Since the NMPC method allows one to compute finite horizon dynamic decision problems, with solutions approximating the corresponding infinite horizon models, it is well suited to track

the solution paths for information constrained agents in the sense of Sims (2005, 2006). Yet, as argued in Section 1 the pitfalls of finite horizon models do not necessarily arise here since we are proposing a solution method using a receding horizon strategy, where one is always $N - 1$ steps away from the endpoint. The NMPC method also permits us to compute state and control variables for models with multiple equilibria, regime changes, and models where the state and control paths needed to be tracked. We can compute discrete and continuous time models where the steady states, as terminal conditions, and linearization about them, are not needed to compute the solutions numerically. We have also shown how NMPC can be applied to high dimensional models to find global solutions. In this case, as well as in the context of multiple equilibria models, we could then show that the decision horizon of the agents, and thus their informational constraints, do matter to find the proper solution. Our examples also show that NMPC can address deterministic and stochastic model variants with good accuracy as well as models with multiple domains of attraction. So far algorithms and software are available that operate in both MATLAB as well as in C++. As compared to Dynamic Programming the NMPC approach, by avoiding to grid the state space, has significant advantages as it is less prone to the curse of dimensionality.

A Parameters for the 5d model of Section 6

Parameter	Definition	Values for simulation
ρ	discount rate	0.03
n	rate of population growth	0.015
α_1	share of public resources used to build up new public capital	0.1
α_2	share of public resources used for the transfers and public consumption	0.7
α_3	share of public resources used for the functioning of the administration	0.1
η	exponent in the utility function	0.1
\bar{M}	pre-industrial level of GHG concentration	1
ε	exponent in the utility function	1.1
ν_1	fraction of public capital to support market activity	0.6 ($\in [0, 1]$)
ν_2	fraction of public capital to mitigate climate change damages	0.3 ($\in [0, 1]$)
ν_3	fraction of public capital to reduce GHG emissions	0.1 ($\in [0, 1]$)
ω	exponent in the utility function	0.05
σ	exponent in the utility function	1.1
A	prductivity of output production	1
A_n	efficiency index of private capital in production	1
A_r	efficiency index of non-renewable resources in production	30 ($\in [10, 100]$)
α	exponent of private capital and resource in the production function	0.5
β	exponent of public capital in the production function	0.5
δ_k	depreciation of private capital	0.075
δ_g	depreciation of public capital	0.05
ψ	parameter in extraction cost function	1
τ	parameter in extraction cost function	2
i_p^f	foreign aid per period earmarked for investment in public capital	0.05
i_p^b	net borrowing per period earmarked for investment in public capital	0
γ	fraction of GHG emissions not absorbed by the ocean	0.9
μ	inverse of the atmospheric lifetime of GHG emissions	0.01
κ	parameter determining the level where GHG cocentration can be stabilized	2
θ	efficiency of public sector's efforts to reduce GHGs	0.01

References

- Amrit, R., Rawlings, J. B. and Angeli, D. (2011). Economic optimization using model predictive control with a terminal cost, *Annual Rev. Control* **35**: 178–186.
- Angeli, D., Amrit, R. and Rawlings, J. B. (2009). Receding horizon cost optimization for overly constrained nonlinear plants, *Proceedings of the 48th IEEE Conference on Decision and Control – CDC 2009*, Shanghai, China, pp. 7972–7977.
- Angeli, D. and Rawlings, J. B. (2010). Receding horizon cost optimization and control for nonlinear plants, *Proceedings of the 8th IFAC Symposium on Nonlinear Control Systems – NOLCOS 2010*, Bologna, Italy, pp. 1217–1223.
- Aruoba, S. B., Fernandez-Villaverde, J. and Rubio-Ramirez, J. F. (2006). Comparing solution methods for dynamic equilibrium economies, *J. Econ. Dyn. Control* **30**: 2477–2508.
- Azariadis, C. and Drazen, A. (1990). Threshold externalities in economic development, *Quarterly Journal of Economics* **105**(2): 501–526.
- Bertsekas, D. P. (2005). *Dynamic Programming and Optimal Control*, Vol. 1, 3 edn, Athena Scientific.
- Bréchet, C., Camacho, C. and Veliov, V. M. (2012). Adaptive model-predictive climate policies in a multi-country setting, Documents de Travail du Centre d’Economie de la Sorbonne 2012.29.
- Brock, W. A. and Mirman, L. (1972). Optimal economic growth and uncertainty: the discounted case, *J. Econ. Theory* **4**: 479–513.
- Brock, W. A. and Scheinkman, J. A. (1976). Global asymptotic stability of optimal control systems with applications to the theory of economic growth, *J. Econom. Theory* **12**(1): 164–190. Hamiltonian dynamics in economics.
- Brock, W. and Starrett, D. (1999). Nonconvexities in ecological management problems, SSRI Working Paper 2026, Department of Economics, University of Wisconsin–Madison.
- Camilli, F. and Falcone, M. (1995). An approximation scheme for the optimal control of diffusion processes, *RAIRO, Modélisation Math. Anal. Numér.* **29**: 97–122.
- Cannon, M., Kouvaritakis, B. and Wu, X. (2009). Probabilistic constrained MPC for multiplicative and additive stochastic uncertainty, *IEEE Trans. Automat. Control* **54**(7): 1626–1632.
- Capuzzo Dolcetta, I. (1983). On a discrete approximation of the Hamilton–Jacobi equation of dynamic programming, *Appl. Math. Optim.* **10**: 367–377.
- Collard, F. and Juillard, M. (2001). Accuracy of stochastic perturbation methods: The case of asset pricing models, *J. Econ. Dyn. Control* **25**: 979–999.
- Couchman, P. D., Cannon, M. and Kouvaritakis, B. (2006). Stochastic MPC with inequality stability constraints, *Automatica* **42**(12): 2169–2174.

- Damm, T., Grüne, L., Stieler, M. and Worthmann, K. (2014). An exponential turnpike theorem for dissipative discrete time optimal control problems, *SIAM Journal on Control and Optimization* **52**(3): 1935–1957.
- Diehl, M., Amrit, R. and Rawlings, J. B. (2011). A Lyapunov function for economic optimizing model predictive control, *IEEE Trans. Autom. Control* **56**: 703–707.
- Dorfman, R., Samuelson, P. A. and Solow, R. M. (1987). *Linear Programming and Economic Analysis*, Dover Publications, New York. Reprint of the 1958 original.
- Falcone, M. (1987). A numerical approach to the infinite horizon problem of deterministic control theory, *Appl. Math. Optim.* **15**: 1–13. *Corrigenda*, *ibid.*, 23 (1991), 213–214.
- Farmer, R., Waggoner, D. and Zha, T. (2009). Understanding Markov-switching rational expectations models, *Journal of Economic Theory* **144**: 1849–1867.
- Feichtinger, G., Kort, P., Hartl, R. F. and Wirl, F. (2001). The dynamics of a simple relative adjustment–cost framework, *German Economic Review* **2**(3): 255–268.
- Greiner, A., Grüne, L. and Semmler, W. (2014). Economic growth and the transition from non-renewable to renewable energy, *Environment and Development Economics*. To appear.
- Greiner, A. and Semmler, W. (2008). *The Global Environment, Natural Resources, and Economic Growth*, Oxford University Press, Oxford, UK.
- Greiner, A., Semmler, W., Diallo, B., Rezai, A. and Rajaram, A. (2007). Fiscal policy, public expenditure composition, and growth. Theory and Empirics, World Bank Policy Research Paper no. 4405.
- Grüne, L. (1997). An adaptive grid scheme for the discrete Hamilton-Jacobi-Bellman equation, *Numer. Math.* **75**(3): 319–337.
- Grüne, L. (2013). Economic receding horizon control without terminal constraints, *Automatica* **49**: 725–734.
- Grüne, L. and Pannek, J. (2011). *Nonlinear Model Predictive Control. Theory and Algorithms*, Springer-Verlag, London.
- Grüne, L. and Semmler, W. (2004). Using dynamic programming with adaptive grid scheme for optimal control problems in economics, *J. Econ. Dyn. Control* **28**: 2427–2456.
- Grüne, L. and Stieler, M. (2014). Asymptotic stability and transient optimality of economic MPC without terminal conditions, *Journal of Process Control* **24**(8): 1187–1196.
- Haunschmied, J. L., Kort, P. M., Hartl, R. F. and Feichtinger, G. (2003). A DNS-curve in a two-state capital accumulation model: a numerical analysis, *J. Econ. Dyn. Control* **27**(4): 701–716.
- Judd, K. and Guu, S.-M. (1997). Asymptotic methods for aggregate growth models, *J. Econ. Dyn. Control* **21**: 1025–1042.

- Judd, K. L. (1998). *Numerical methods in economics*, MIT Press, Cambridge, MA.
- Juillard, M. and Villemot, S. (2011). Multi-country real business cycle models: Accuracy tests and test bench, *Journal of Economic Dynamics and Control* **35**(2): 178 – 185.
- Kaganovich, M. (1985). Efficiency of sliding plans in a linear model with time-dependent technology, *The Review of Economic Studies* **52**(4): 691–702.
- Kerrigan, E. C. (2000). Robust constraint satisfaction: Invariant sets and predictive control, PhD Thesis, University of Cambridge.
- McKenzie, L. W. (1986). Optimal economic growth, turnpike theorems and comparative dynamics, *Handbook of Mathematical Economics, Vol. III*, North-Holland, Amsterdam, pp. 1281–1355.
- Parra-Alvarez, J. C. (2012). A comparison of numerical methods for the solution of continuous- time DSGE models, Manuscript, University of Aarhus.
- Rawlings, J. B. and Mayne, D. Q. (2009). *Model Predictive Control: Theory and Design*, Nob Hill Publishing, Madison.
- Santos, M. S. and Vigo-Aguiar, J. (1995). Accuracy estimates for a numerical approach to stochastic growth models, Discussion Paper 107, Institute for Empirical Macroeconomics, Federal Reserve Bank of Minneapolis.
- Santos, M. S. and Vigo-Aguiar, J. (1998). Analysis of a numerical dynamic programming algorithm applied to economic models, *Econometrica* **66**: 409–426.
- Scheinkman, J. (1976). On optimal steady states of n -sector growth models when utility is discounted, *Journal of Economic Theory* **12**: 11–30.
- Sims, C. A. (2005). Rational inattention: a research agenda, Discussion Paper Series 1: Economic Studies, 34, Deutsche Bundesbank Research Centre.
- Sims, C. A. (2006). Rational inattention: Beyond the linear-quadratic case, *American Economic Review, American Economic Association* **96**: 158–163.
- Skiba, A. (1978). Optimal growth with a convex-concave a production function, *Econometrica* **46**(3): 527–539.
- von Neumann, J. (1945). A model of general economic equilibrium, *The Review of Economic Studies* **13**(1): 1–9.
- Zaslavski, A. (2006). *Turnpike Properties in the Calculus of Variations and Optimal Control*, Springer, New York.
- Zaslavski, A. (2014). *Turnpike Phenomenon and Infinite Horizon Optimal Control*, Springer International Publishing.