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

February 21, 2013

Lars Grüne,[1] Willi Semmler[2] and Marleen Stieler[3]

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 moderate than for DP. In this paper we explain the basic idea of NMPC together with some implementational details and illustrate its ability to solve dynamic decision problems in economics by means of numerical simulations for various examples, including stochastic problems, models with multiple equilibria and regime switches in the dynamics.

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

*This paper has been started when Lars Grüne visited the University of Technology, Vienna, and Willi Semmler was Fulbright Professor in Vienna. Willi Semmler would like to thank the Fulbright Commission for the research support. We want to thank Vladimir Veliov and the research group on control theory and nonlinear dynamics at the University of Technology, Vienna. The paper has been presented at the 6th international conference on Computation in Finance and Economics, Oviedo, Spain. Comments from the participants and Michel Juillard are gratefully appreciated. Lars Grüne was supported by the European Union under the 7th Framework Programme 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.

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

[2]New School For Social Research, New York, email: semmlerw@newschool.edu

[3]Mathematical 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 is provided in Santos and Vigo-Aguiar (1998), Judd (1998), and Grüne and Semmler (2004).[4] The latter have proposed Dynamic Programming (DP), with grid refinement, 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. As other numerical procedures, applied to solve dynamic decision models, such as provided by some numerical methods to solve DSGE models, DP makes strong assumption on the knowledge and information that agents should be endowed with in the context of an infinite horizon model.

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). Similarly to DP, NMPC can solve nonlinear dynamic decision problems 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.

The use of NMPC in an economic context has two different interpretations. On the one hand, in the references cited above it was shown that, under suitable assumptions, for sufficiently long finite optimization horizon and/or by employing judiciously chosen constraints, NMPC yields approximately infinite horizon optimal trajectories. On the other hand, due to the use of finite horizon optimal trajectories, NMPC also allows to study decision making in a shorter time horizon or agents making decisions over different time horizons. Hence, compared to the infinite horizon models to be solved with their strong informational requirements, there is much less requirement of information that agents need to have when making decisions on a finite horizon in the NMPC context. As such, the NMPC approach can be seen as a particular way to implement the concept of rational inattention in dynamic

[4] In Grüne (1997) and Grüne and Semmler (2004) an adaptive grid scheme is used for finding global solutions of models with dynamic optimization. Those numerical methods provide us with approximate solutions and accuracy estimates for the numerical methods can be employed.

decision making.

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 continuously available information. In this context, we can interpret the NMPC solutions as based on agents' decision making using only limited information. Yet, if the agents information and information processing capacity increases this is likely to approximate better the infinite horizon decision making[5], which is reflected, e.g., in the examples in Sections 3.1 and 4.2, below. Moreover, NMPC permits to study models with parameter uncertainties and learning,[6] models with multiple equilibria, and jumps and regime changes in the dynamics. One also can easily track the dynamics of the state variables over a finite horizon and paths with regime switching. NMPC does not require to know the terminal or steady state conditions, yet, if such information is available the performance can be improved by incorporating it into the algorithm as terminal constraints.

The goal of this paper is to evaluate the performance of NMPC via computer simulations for a selection of dynamic decision models in economics. We consider discounted continuous time and discrete time optimal decision problems of one to three dimensional economic models, with one to three control variables. In order to study the ability of NMPC to approximate the optimal infinite horizon solution, 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, and 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.

There are, however, in the economic literature more complicated dynamic models with optimizing agents which have been a challenge to commonly used numerical techniques. Those models exhibit more complicated dynamics due to the existence of multiple domains of attraction,[7] thresholds and regime changes, parameter uncertainty and learning etc. Examples of such models can be found in the literature on economic growth,[8] macrodynamics of RBC and DSGE type, in growth models with exhaustible resources, as well as in dynamic models with the financial sector and credit markets. Multiple steady states and thresholds can also arise in the dynamic decision problem of the firm, in resource economics and in ecological management problems.[9] 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. Section 3 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 and

[5]Sims 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)

[6]See Bréchet et al. (2012)

[7]If there are local attractors for some equilibria authors characterize them as indeterminant equilibria. For a survey of models with indeterminacy, see Benhabib and Farmer (1999).

[8]In 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).

[9]For the former see, Feichtinger et al. (2001) and Hartl et al. (2000), for the latter see in particular Brock and Starrett (1999)

some extension of it. In Section 4 we then study two dimensional dynamic optimization models, a stochastic version of the basic growth model, a model with multiple domains of attraction and threshold, and another one where the paths of the state variables need to be tracked. Section 5 then presents results for three dimensional models both of the exhaustible resource type where also the solution paths need to be tracked as well as a credit market model with regime changes. Section 6 gives details of the numerical implementation and discusses some known pitfalls of the NMPC method. Section 7 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. We consider infinite horizon discounted optimal control problems, either given in continuous time $t \in \mathbb{R}_0^+$ by

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

where

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

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

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

where

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

and $\mathcal{U}$ and $\mathcal{U}_d$ are appropriate sets of control functions and control sequences, respectively.

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 method[10]. To this end, the continuous time optimal control problem (2.1)–(2.2) is replaced by a first order discrete time approximation given by

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

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

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

[10]This 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).

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 $\tilde{x}(k)$.

2.1 The idea of NMPC

The idea of NMPC now lies in replacing the maximization of the discrete time infinite horizon functional (2.3) by the iterative maximization of finite horizon functionals

$$\max_{u_i \in \mathcal{U}} \sum_{k=0}^{N-1} \beta^k g(x_i(k), u_i(k)), \tag{2.7}$$

for a truncated finite horizon $N \in \mathbb{N}$ with given initial value $x_i(0) \in \mathbb{R}^n$ and x_i generated by the usual dynamics $x_i(k+1) = \varphi(x_i(k), u_i(k))$ for $k = 0, 1, 2, \ldots, N-1$. The index i indicates the number of the iteration. Note that in the standard case, neither β nor g nor φ changes when passing from (2.3) to (2.7), only the optimization horizon is truncated, though changes can be allowed for in extensions, see below and Bréchet et al. (2012).

Problems of type (2.7) can be efficiently solved numerically by converting them into a static nonlinear program and solving them by efficient NLP solvers, implementation details are discussed in Section 6, below.

Given an initial value $x(0)$, NMPC now generates solutions on an infinite time horizon by iteratively solving (2.7) as follows:

(1) for $i = 0, 1, 2, 3, \ldots$
(2) solve (2.7) with initial value $x_i(0) := x(i)$ and denote the resulting optimal control sequence by $u_i^*(\cdot)$
(3) set $u(i) := u_i^*(0)$ and $x(i+1) := \varphi(x(i), u(i))$
(4) end of for-loop

This algorithm yields an infinite trajectory $x(i)$, $i = 0, 1, 2, 3, \ldots$ whose control sequence $u(i)$ consists of all the first elements $u_i^*(0)$ of the optimal control sequences for the finite horizon subproblems (2.7). In what follows, we refer to the finite horizon optimal trajectories corresponding to $u_i^*(\cdot)$ computed in Step (2) as the *open loop* trajectories while the trajectory $x(i)$ computed in Step (3) will be referred to as the *closed loop* trajectory.

Under appropriate assumptions on the problem, it can be shown that the closed loop solution $(x(\cdot), u(\cdot))$ (which depends on the choice of N in (2.7)) approximates the optimal solution of (2.3) if the time horizon N is sufficiently large. The reason for this at a first glance surprising behaviour relies on the so called turnpike property, cf. McKenzie (1986). This property ensures that the finite horizon optimal open loop trajectories pass through neighbourhoods $B_\varepsilon(x^\star)$ of an infinite horizon optimal equilibrium $x^\star$ with $\varepsilon \to 0$ as $N \to \infty$ (for an illustration see Figure 6.12). This behaviour implies that the piece of the finite horizon optimal trajectory connecting the initial value $x_i(0)$ to $B_\varepsilon(x^\star)$ is approximately optimal among all trajectories connecting $x_i(0)$ to $B_\varepsilon(x^\star)$ and hence also approximately optimal for the infinite horizon problem since the infinite horizon trajectory will also enter

$B_\varepsilon(x^\star)$. Consequently, since the NMPC closed loop solution is synthesized from these approximately optimal trajectory pieces, we can also expect it to be approximately optimal. These arguments are made mathematically precise in Grüne (2013). While the proof in this reference is carried out for averaged optimal control problems, due to the fact that the ingredients of the proof are also known to hold for discounted problems, we expect the same property to hold for the discounted setting considered here which is confirmed by the numerical example in Section 3.1. In any case, the solution generated by NMPC will always provide a lower bound for the infinite horizon optimal solution.

As already mentioned in the introduction, in case N is not large enough such that the NMPC solution approximates the infinite horizon optimal trajectory, the solution still has a valid economic interpretation as an approximate solution based on decisions taken on the basis of incomplete information. Here, incomplete refers to the fact that only the future behaviour on a finite horizon is taken into account in decision making.

Even though (2.7) can typically be efficiently solved numerically, the quality of the solution strongly depends on the length of the horizon N and for large N sometimes numerical problems can be observed and the computed solution deteriorates[11]. If an approximation of the infinite horizon optimal solution is desired but the problem (2.3) is numerically infeasible for the needed horizon lengths N, then one can improve the performance by using the infinite horizon optimal steady state $x^\star$ 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 $u_{k,i}$ for which $x_{N,i} = x^\star$ holds. Again, the benefit of adding terminal constraints is so far only rigorously proven for averaged optimal control but the example in Section 3.2 shows that the positive effect of adding this equilibrium terminal constraint is also visible in the discounted setting. We note that such terminal constraints were only used in this example; all other computations in this paper were performed without using terminal (steady state) constraints when solving (2.3).

2.2 NMPC for stochastic problems

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

$$V(x_0) = E\left(\max_{u\in\mathcal{U}} \sum_{k=0}^{\infty} \beta^k g(x(k), u(k))\right) \tag{2.8}$$

with the discrete time stochastic dynamics

$$x(k+1) = \varphi(x(k), u(k), z_k), \quad x(0) = x_0 \in \mathbb{R}^n, \tag{2.9}$$

where the z_k are i.i.d. random variables. Again, this problem could be a priori in discrete time or it could be 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).

[11] It strongly depends on the model for which N these numerical problems become visible. In our examples, these typically occurred for values around $N \approx 80$–100.

From a computational point of view, the main difficulty in stochastic NMPC is the efficient solution of the corresponding finite horizon problem (2.7) 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 which does in general not compute the true stochastic optimum but in the case of stochastic perturbations with low intensities typically still yields reasonably good approximately optimal results. To this end, we replace the stochastic dynamics by its expected counterpart

$$x^e(k+1) = E\Big(\varphi(x^e(k), u(k), z_k)\Big), \quad x^e(0) = x_0 \in \mathbb{R}^n \tag{2.10}$$

and in each iteration instead of (2.7) we solve

$$\max_{u_i \in \mathcal{U}} \sum_{k=0}^{N-1} \beta^k g(x_i^e(k), u_i(k)). \tag{2.11}$$

Note that we only use (2.10) in order to solve (2.11) in Step (2) of the NMPC algorithm. In Step (3) we simulate the closed loop using the original stochastic dynamics (2.9) with z_k realized by appropriate random numbers.

3 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.

3.1 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 literature. 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

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

with constants $A > 0$ and $0 < \alpha < 1$. The exact solution to this problem is known (see Santos and Vigo-Aguiar (1998)) and is given by

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

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

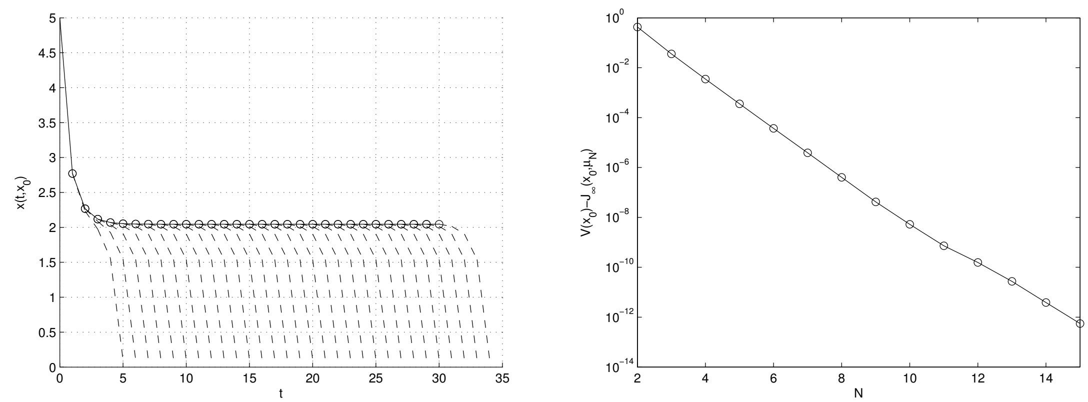

Figure 3.1: Closed loop trajectory (solid) and open loop trajectories (dashed) for the growth model for $N = 5$ and $x_0 = 5$ (left) and $V(x_0) - J_\infty(x_0, \mu_N)$ for $x_0 = 5$, $N = 2, \ldots, 15$ (right), with μ_N the NMPC controller for time horizon N.

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 3.1 (left; the closed loop solution is depicted as solid line with circles). Moreover, we see that the open loop trajectories (shown as dashed lines) converge into a neighborhood of $x^\star$, stay there for some time instants and finally turn away. Numerical tests show, that the time period the open loop trajectories stay nearby the equilibrium increases as we enlarge the optimization horizon N. This phenomenon, often referred to as *turnpike property*, can be observed for several types of optimal control problems (see also Figure 6.12)) and is an essential ingredient in order to prove convergence of the closed loop trajectory for averaged optimal control problems Damm et al. (2012).

The second interesting aspect to look at is whether the NMPC closed loop trajectory maximizes the given objective function. In Figure 3.1 (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 3.1 shows, that the difference $V(x_0) - J_\infty(x_0, \mu_N)$ converges to zero exponentially fast for $N \to \infty$ and hence, for increasing N, the NMPC generated return $J_\infty(x_0, \mu_N)$ approximates the optimal value $V(x_0)$ arbitrarily well.

3.2 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 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.[12] 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

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

with dynamics

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

with u_1 consumption, u_2 labor effort, x as capital stock, $r(t) = \alpha A x(t)^{1-\alpha} u_2(t)^{1-\alpha}$ 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 $A = 1$, $\alpha = 0.4$, $\delta = 0.0196$, $\sigma = 1.8011$. The discount rate for the dynamic decision problem is taken as $\rho = 0.0105$. The steady state values are[13] $x^* = 23.03$, $u(1)^* = 1.2865$, $u(2)^* = 0.3104$.

In Figure 3.2 the closed loop solution and open loop solution are shown, using the NMPC algorithm, in both cases the steady state of the capital stock, $x^* = 23.03$, is used as terminal condition.

The above example of Section 3.1 is in fact computed without terminal condition. Also the next examples are computed without terminal conditions. In any case, if avoidable, one should not use linearization techniques about some known steady states, but rather find global solutions, in particular if some interesting dynamics further away from the steady state. This is even more important if there are multiple domains of attraction, as our next example shows.

[12] The 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. Recently, 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).

[13] For details see Parra-Alvarez (2012)

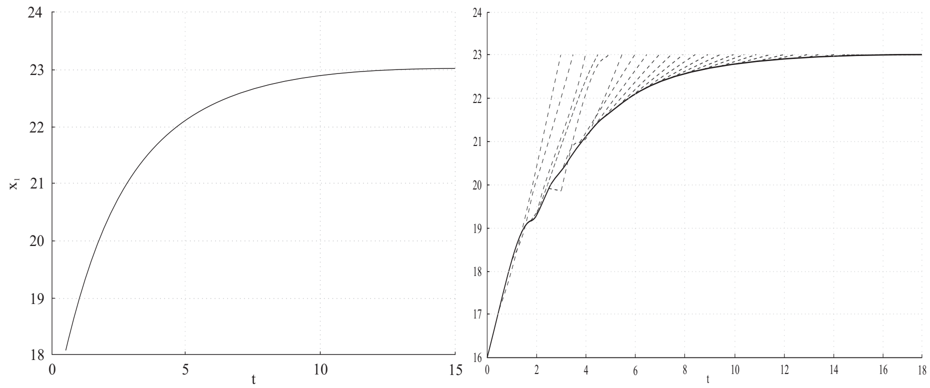

Figure 3.2: Closed loop solution only (left) and closed loop and open loop solutions (right), in both cases steady state of capital stock $x^* = 23.03$ used as terminal condition.

4 Two dimensional examples

Subsequently, we describe three two–dimensional problems which we have solved using our algorithm. We first study our above basic growth model of Sect. 3.1, but we employ an extension to a two dimensional stochastic variant by which means we illustrate the capability of NMPC for solving stochastic problems.

We then turn to deterministic problems, where the study of the dynamics of two–dimensional problems with multiple domains of attractions 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 interesting to track the finite time path of the solution trajectories.

4.1 A 2d stochastic growth model

Our first 2d problem is a 2d version of the Brock and Mirman (1972) model of Example 3.1. Here the 1d model from Example 3.1 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 $x_1(k+1) = \alpha\beta A x_2(k) x_1(k)^\alpha$. 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 4.3 (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 $\widetilde{E}(x_1(k))$ 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 4.3 (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.

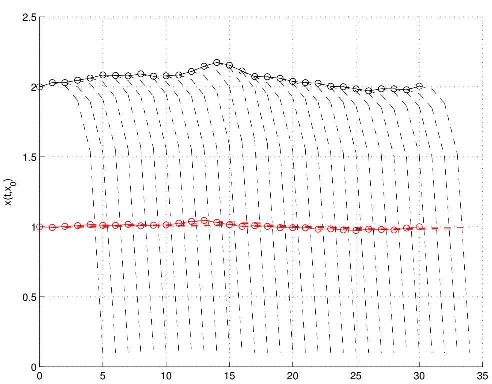

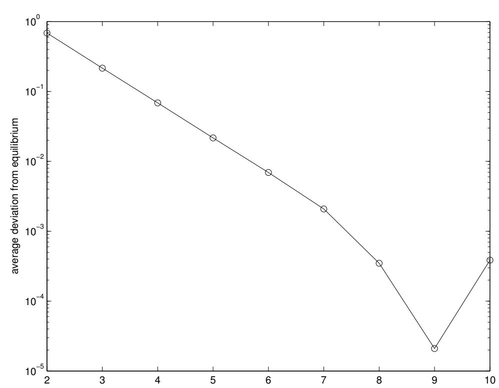

Figure 4.3: Closed loop trajectory (solid) and open loop trajectories (dashed) for the growth model for $N = 5$ and $x_0 = (2,1)^T$ (left) and equilibrium deviation $|\widetilde{E}(x_1(t)) - x_1^{e,*}|$ for $x_0 = x^{e,*}$, $N = 3, \ldots, 15$ (right), with μ_N the NMPC controller for time horizon N.

4.2 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

$$g(x,u) = k_1\sqrt{x_1} - \frac{x_1}{1 + k_2 x_1^4} - c_1 x_2 - \frac{c_2 x_2^2}{2} - \frac{\alpha u^2}{2}$$

with the dynamics

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

The following parameters are used: $\sigma = 0.25$, $k_1 = 2$, $k_2 = 0.0117$, $c_1 = 0.75$, $c_2 = 2.5$, $\alpha = 12$ and discount rate $\delta = 0.04$. Generating the vector field we have obtained the results shown in the following Figure 4.4.[14]

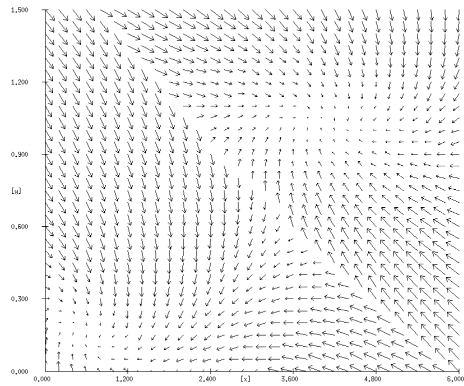

Figure 4.4: Vector field of the model with multiple attractors

The vector field shows clearly two domains of attraction, 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. Note, however, that the above vector field is generated by DP[15] and assumes an infinite decision horizon.

Next, we want to see 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 4.5 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 4.4, 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. So, if one views the above model as a data generating mechanism, data would be generated on different time scales, as wavelet approaches suggest, see Gallegati et al. (2011).

4.3 A 2d growth model with non-renewable resources

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

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

[14]Figure 4.4 was generated through DP, with time step $h = 1/20$ and 101 control values $\widetilde{U}$ covering $U = [-1, 1]$. For the details of solving the model through DP and generating the vector field, see Grüne and Semmler (2004).

[15]See Grüne and Semmler (2004)

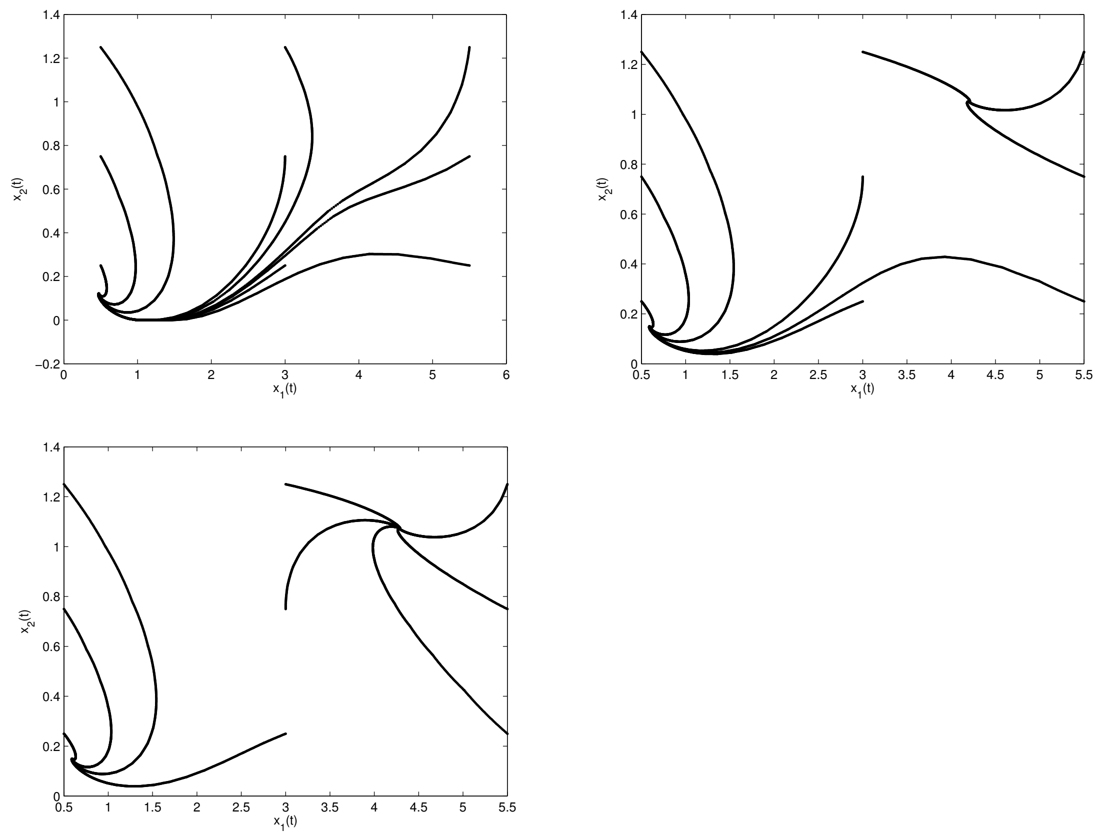

Figure 4.5: Phase plots of NMPC trajectories of the model of multiple domains of attraction for $N = 10$ (upper left), $N = 35$ (upper right) and $N = 50$ (lower)

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

$$x_2(t) \leq x_2(0)$$

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 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 $x_2(t) \leq x_2(0)$ 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 4.6 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

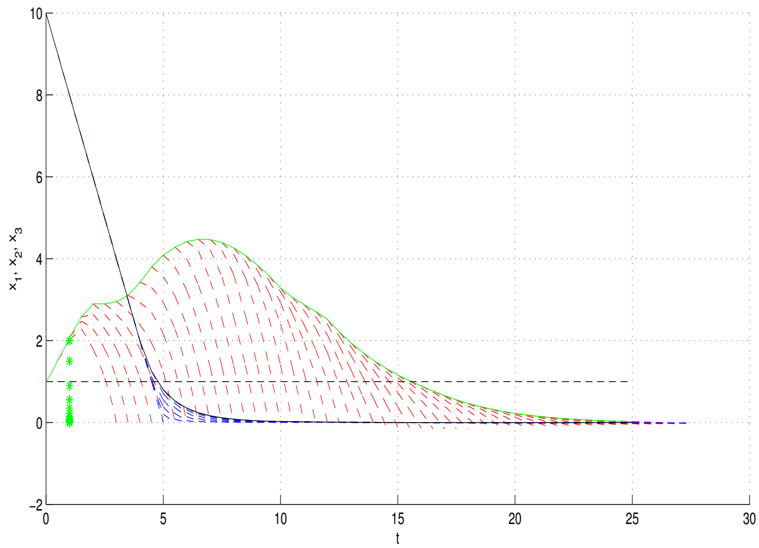

Figure 4.6: Growth model with non-renewable resource, closed loop solid, open loop dashed, initial value $x(0) = (1, 10)^T$

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.[16] 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 study 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 we introduce a backstop technology as is done next.

5 Three dimensional examples

Next we are discussing three dimensional models that have three state variables and two and three decision variables respectively. We start with a model of two decision variables.

5.1 Growth with non-renewable resources and backstop technology

We want to extend the model of Section 4.3 by considering the time paths of the state and decision variables when externalities arise in the context of a growth model with non-renewable resources (such as fossil fuel, with CO_2 emission) are used. We will allow, however, a backstop technology that might be phased in.[17] The issue here is then whether

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

[17] This type of model originates in the work of Heal. For a recent review of the concept of backstop technology, see for example Heal (2009). For further literature references and modeling details, see Greiner et al. (2012).

the non-renewable resource is completely exhausted, as in the case of Sect. 4.3, or whether some is left in situ when the backstop technology is phased in.

In the subsequent model, which is based on the above literature, there are two control variables and three state variables. The preferences include a utility from consumption, u_1, affected multiplicatively by the cumulative CO_2 emission, x_3, arising from the use of non-renewable energy such as fossil fuel. Preferences are written as:

$$U = \frac{u_1^{1-\sigma}(x_3 - x_o)^{-\xi(1-\sigma)} - 1}{1-\sigma} \tag{5.12}$$

Here again, the parameter $1/\sigma > 0$ denotes the inter-temporal elasticity of substitution of consumption between two points in time and $\xi > 0$ gives the (dis)utility of the greenhouse gas concentration, exceeding the pre-industrial level, x_o, affecting negatively consumption. For $\sigma = 1$ the utility function is logarithmic in consumption and pollution.

The dynamics of the three state variables, capital stock, stock of non-renewable resource and stock of CO_2, x_1, x_2, and x_3 are as follows:

$$\dot{x}_1(t) = Y(t) - u_1(t) - \delta x_1(t) - a \cdot u_2(t) \tag{5.13}$$
$$\dot{x}_2(t) = -u_2(t) \tag{5.14}$$
$$\dot{x}_3(t) = \beta_1 u_2(t) - \mu(x_3(t) - \kappa x_o) \tag{5.15}$$
$$x_2(0) = x_0, x_1(0) = x_0, x_3(0) = x_0$$

We have $u_2(t)$ as the extracted amount of non-renewable energy (fossil fuels) used at time t to generate energy and x_1 a stock of capital that produces energy using renewable sources of energy (backstop technology, such as wind or solar energy). Note that production of the final good $Y(t)$ uses energy and is a concave function of energy input,[18]

$$Y = A\,E^{\alpha} = A\,(A_n x_1 + A_p u_2)^{\alpha} \tag{5.16}$$

and A_i, $i = p, n$, denote efficiency indices. Total energy E consists of the sum of these two types of energy, non-renewable and renewable,with $0 < \alpha \leq 1$, $A > 0$. Note that energy is a homogeneous good so that modeling the two types as perfect substitutes can be justified.

In the numerical analysis below we set $\mu = 0.1$ and $\beta_1 = 0.5$ which are plausible values and $\kappa = 2$, $M(0) = M_o$ with $M_o = 1$. Furthermore, the other parameters are $\rho = 0.05$, $\delta = 0.05$, $\sigma = 1$, $\alpha = 0.5$, $A_p = 1$, $A_n = 1$, $a = 0.1$, and $\xi = 0.5$.

Figure 5.7 shows the NMPC results for different initial conditions of the non-renewal resource, $x_2(0) = 0.5$, $x_2(0) = 2.5$. In both cases a high initial capital $x_1(0) = 3$ exists and also in both cases we assume $x_3(0) = 1$.

[18] In the following we delete the time argument t as long as no ambiguity arises.

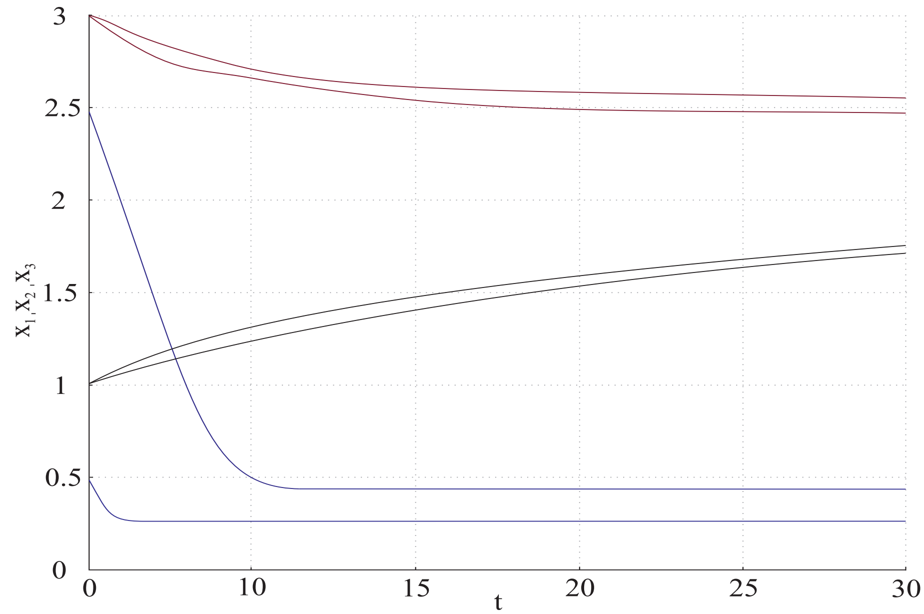

Figure 5.7: NMPC closed loop solutions for $x_2(0) = 0.5$, and $x_2(0) = 2.5$, and in both cases $x_1(0) = 3, x_3(0) = 1$

In Figure 5.7, for both the low $x_2(0) = 0.5$ and high $x_2(0) = 2.5$, there is non-renewable energy unextracted if the initial capital stock is high and thus sufficient renewable energy produced right away. As our simulations have shown[19], for both, $x_2(0) = 0.5$, and $x_2(0) = 2.5$, the cumulative CO_2 is lower when the renewable energy, the backstop technology, is phased in faster, i.e. when the capital stock is high. We also can observe that for both initial conditions of the non-renewable resource, some of the non-renewable resource is left unextracted. Yet, for the second case, for $x_2(0) = 0.5$, the cumulative CO_2 is lower.

Thus, in contrast to what the Hotelling theorem seems to imply[20], with backstop technology available, not all the resource will be extracted, non-renewable energy resource will be left in situ, a quicker phasing in of renewable energy will take place, and, as compared to the complete extraction of the resource, less CO_2 is emitted. If the available non-renewable energy resource is high, the extraction rate of the non-renewable resource is initially high, this will keep the CO_2 emission high, and we end up with a high concentration of CO_2 sooner.[21] So, NMPC, by tracking the solution paths, helps to provide us with useful information on the fate of the state variables.

[19] For details, see Greiner et al. (2012).

[20] The conjecture that the Hotelling model of the optimal use of a non-renewable resource, where only one energy source is available, will lead to a faster depletion, to a faster build up of CO_2 emission and faster global warming, depends of course on the set-up of the Hotelling model and how prices and costs are modeled. For results on this conjecture in a Hotelling model using firms' payoffs, see Maurer and Semmler (2010) and for a model with preferences see section 4.3.

[21] Note that also in Figure 5.7, with high initial non-renewable energy, more fossil energy source is left unextracted, but in this case the cumulative CO_2 rises faster.

5.2 A model with credit market frictions

The next example is a model with credit market frictions. Numerous model variants along this line of research have been proposed. Here, regime changes can occur in such models when suddenly, due to overborrowing, credit constraints or a jump in credit spreads arise.[22]

A basic model with credit market frictions can be described as follows. There are three state variables, employment, x_1, capital, x_2, and debt, x_3. The model includes search and matching frictions in the labor market and in the credit market, with preferences over consumption and leisure given by:

$$U = \frac{u_1^{1-\eta}}{1-\eta} - e x_1^{\chi}$$

$$\dot{x}_1(t) = m^L(su(t), u_2(t)) - \sigma x_1(t) \quad (5.17)$$

$$\dot{x}_2(t) = m^B(u_3(t), B(t)) - \delta x_2(t) \quad (5.18)$$

$$\dot{x}_3(t) \quad = r x_3(t) - \upsilon[Y\left(x_1(t), x_2(t)\right) - F(t) - \Phi\left(s\right)\left(1 - x_1(t)\right) - \zeta u_2(t) - \kappa B(t)] \quad (5.19)$$

with u_1, aggregate consumption, Y, aggregate production, generated from a Cobb-Douglas production function $Y\left(x_1(t), x_2(t)\right) = x_2(t)^{0.36} A x_1(t)^{0.64}$, A, (exogenous) labor productivity, $F(t)$, financial funds, $B(t) = F(t) - u_1(t)$, offered bonds for firms, x_3, the stock of external debt, $u = (1 - x_1)$, the unemployment rate, u_2, the job vacancies posted, r, the interest rate, and x_1, employment. In addition to costly search on the labor market, $\zeta u_2(t)$, there is floating cost for bonds, with per-period bond cost floating, $\kappa B(t)$.

In equs. (5.17) and (5.18) the function $m^L(u(t), u_2(t)) = (su(t))^{0.5} u_3(t)^{0.5}$ denotes the search and matching function on the labor market, with s the search intensity of the unemployed, and $m^B(u_3(t), B(t)) = u_3(t)^{0.5} B(t)^{0.5}$ is a similar function for the credit market. Both are Cobb-Douglas, with exponents 0.5 and 0.5 and σ, the separation rate, and δ, the depreciation rate of capital. In equ. (5.19), x_3 denotes external debt, the term [.] is external borrowing, used for excess spending over domestic income for example by households and firms.[23] The debt dynamics (5.18) is written here in a way which is standard if one allows for external borrowing with no credit constraints, see Blanchard and Fischer (1989, Ch. 2). What is only added here are the search and matching frictions on the labor and credit markets.

Note that we assume $F(t) = \mu u_1(t)$, with $\mu > 0$, so that there is excess of funding over consumption – as long as the latter is optimal – which will be available for domestic investment. This assumption is made in order to avoid a fourth decision variable. Yet, investment funding will be evaluated on the credit market by a search and matching process for bonds of firms. Investment is expressed as intended bond issuing, $B(t)$, but it faces a search and matching process on the credit market, where the supply of funds for bonds is given by

[22] For studies on a sudden rise of credit spreads, see Gilchrist and Zakrajek (2011) and Roch and Uhlig (2012) and the literature referenced there. Both cases, credit constraints and credit spreads are treated in Ernst and Semmler (2012).

[23] Sovereign debt could be included here as well.

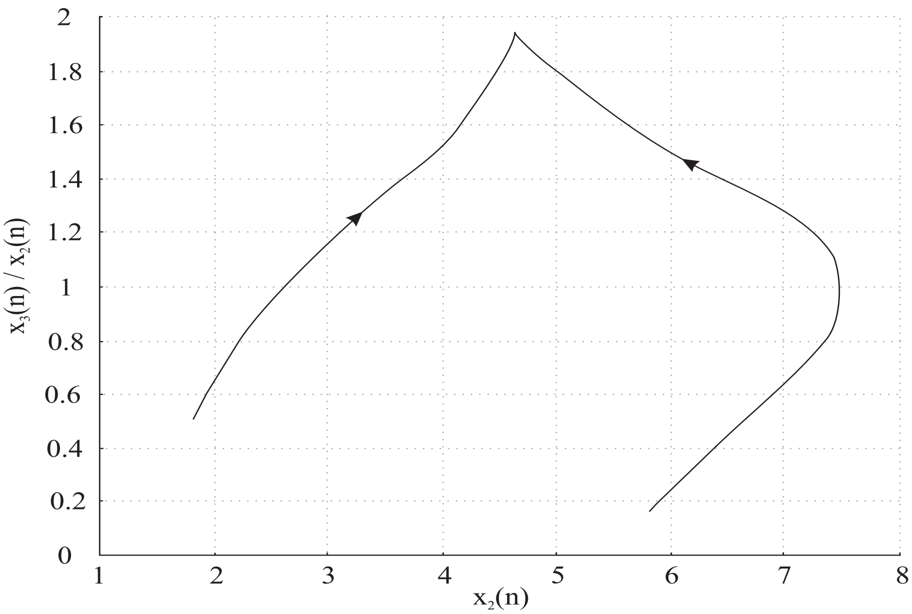

Figure 5.8: Dynamics of a model with no credit constraints, constant interest rate

$F - u_1$. It also means that the screening and monitoring of investment funding takes place more extensively than funding for (optimal) consumption.

The parameters for the NMPC solutions are: $\mu = 1.3$ and $\beta = 0.35$, $\kappa = 0.1$, $\rho = r = 0.03$, $\delta = 0.03$, $\sigma = 0.04$, $\alpha = 0.36$, $A = 1$, $\xi = 0.07$, $\chi = 5$, $e = 1$. The parameter v is set to one in the no-credit constrained regime.

Figure 5.8 shows the results for the basic model. The horizontal axis represents capital stock and the vertical the leveraging ratio, the debt to capital stock. The solution is obtained by NMPC with horizon $N = 6$. Figure 5.8 shows that there is clearly a steady state at about $x_2^* = 4.65$ and $(x_3/x_2)^* = 1.95$. The steady state is unique and all initial conditions for the state variables would converge toward that point.

So far the basic model implies no further credit restrictions. Yet, the credit market could be shut down and no further credit obtained externally if some debt capacity is reached. We could thus have a regime switch such that the term [.] in equ (5.19) becomes zero. This gives rise to a regime change from a smoothly working credit market to a credit constrained capital market,[24] as often formulated in the literature on credit constrained regimes, occurring when the debt constraint holds:

$$x_3(t) \leq \gamma x_1(t) \tag{5.20}$$

where γ is the (exogenously given) debt capacity. When equality starts to hold in (5.20), we have $0 = v[Y(x_2, Ax_1) - F - \Phi(s)(1 - x_1) - \zeta u_2 - \kappa B]$ with $v = 0$ and the economy

[24]Regime switching models in the DSGE literature can be found with respect to regime switches in policy reaction function, technology processes and in nominal rigidities, see Eo (2009) whose work is based on Sims (2002). There, however, it is assumed that the Euler equation, based on an infinite horizon solution, holds.

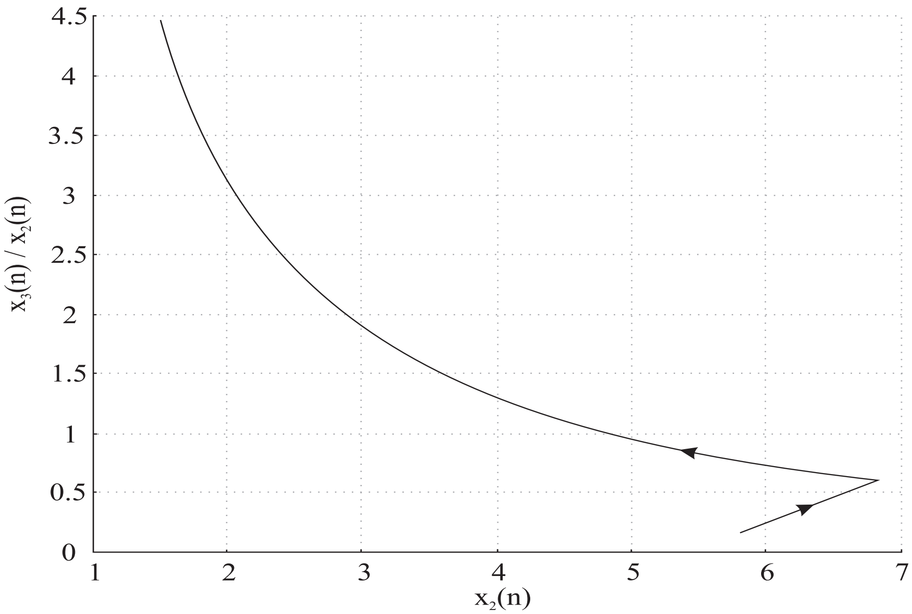

Figure 5.9: Dynamics with regime switches in the credit market

experiences a credit constrained regime,[25] (with likely severe financial stress).[26]

The regime change dynamics, triggered through the credit constraint, can also be explored by our NMPC procedure.[27] Note that consumption continues to be chosen optimally whereby the remainder of output minus consumption (adjusted for the search cost for jobs, and bond issuing) goes into investment. Investment will now be constrained if external funding and credit flows are terminated. The results are shown in Figure 5.9.

Figure 5.9, with the switch to a credit constrained regime, shows that the capital stock as well as debt to capital ratio first go up, then debt rises and finally the capital stock shrinks after the regime switch.

In contrast to the above regime switch, alternatively, as other literature assumes,[28] there could be a higher default premia and credit spread triggered for the borrowers giving rise to the following model variant. In the debt dynamics

$$\dot{x}_3 = r(x_3(t)/x_2(t))x_2(t) - \upsilon[Y_t\left(x_2(t), x_1(t)\right) - F(t) - \Phi\left(s\right)\left(1 - x_1(t)\right) - \zeta u_2(t) - \kappa B(t)] \tag{5.21}$$

[25]Earlier growth literature has referred to such a case in an open economy growth model, see Barro et al. (1995). Of course, some probability of regime switching could be introduced, as in Eo (2009).

[26]In the numerics the regime switch is achieved by first setting parameters before the appropriate terms equal to 1 and then when the constraints set in they are set to 0.

[27]Note that the discontinuities induced into the model by the regime switch may in principle cause problems both in the underlying optimization algorithm as well as in the ODE solver used in the numerical procedure. However, for the parameters used in our simulations no such difficulties were observed, probably because the solutions always cross the regime switching surface transversally.

[28]See Gilchrist and Zakrajek (2011) and Roch and Uhlig (2012)

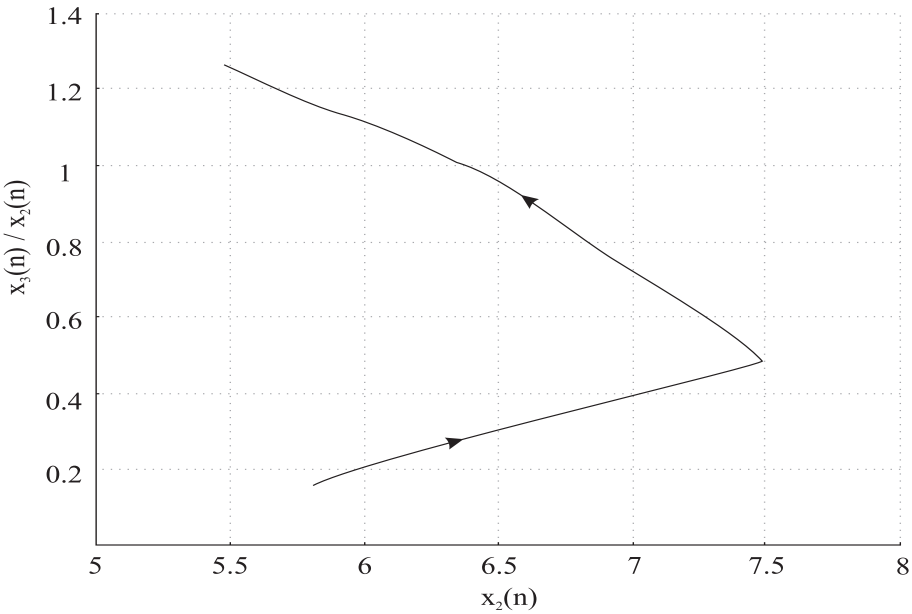

Figure 5.10: Dynamics with regime change in credit spreads

the credit spread is now made dependent on the leverage ratio[29], but we have again $v = 1$.

This formulation is similar to Roch and Uhlig (2012) who have, however, an on-off scenario: With a high probability of default bond prices are low and yields are high, and the reverse holds for a low probability of default. If one wants to smooth out the on-off cases, as the only two scenarios, we can perceive a continuum of cases where the probability of default may steadily rise starting from a low level, then rising faster, and then leveling off, where no bonds can be issued any more. One can thus make the bond price and thus the yield, a nonlinear function of leveraging.[30]

In parallel to Roch and Uhlig (2012) we make the credit spread an arctan function of the leverage, so that the credit spread becomes:

$$r(x_3/x_2) = \beta arctan(x_3/x_2). \tag{5.22}$$

which in turn is likely to give rise to a shrinking consumption and investment demand, reducing output and the surplus, to service the debt.[31]

Figure 5.10 provides the NMPC results when credit spread follows the nonlinear function (22). As can be observed from Figure 5.10, as compared to Figure 5.9, the change to a high credit spread economy, possibly triggering a stage of high financial stress and an economic contraction,[32] occurs less abruptly than in Figure 5.9 where there is a sudden regime switch when the maximum borrowing capacity is reached. Yet, in Figure 5.10 we

[29]For further details, Ernst and Semmler (2012) and also Gilchrist and Zakrajek (2011).

[30]Gilchrist and Zakrajek (2011) have added a persistent shock to the leverage ratio to obtain higher bond yields and thus greater credit spreads.

[31]For a mechanism explaining the further downward spiral, see Ernst and Semmler (2012).

[32]How a self-enforcing mechanism, a vicious cycle, might set in is also discussed in Roch and Uhlig (2012).

can also observe the contractionary effect after the regime change to higher default premia and greater spreads.

6 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 routine[33] (for less demanding examples) or by a C++ software package[34] (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

$$\begin{aligned} &\min_{z\in\mathbb{R}^{n_z}} f(z) \\ &\text{s.t.} \quad G(z)=0 \ \text{ and } \ H(z)\le 0. \end{aligned} \tag{6.23}$$

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 proposed[35]:

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 satisfies the system dynamics. This leads to the following definitions in (6.23):

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

$$G(z) := \begin{pmatrix} *_1 \\ x_{0,i}-x_0 \\ x_{1,i}-\varphi(x_{0,i},u_{0,i}) \\ \vdots \\ x_{N,i}-\varphi(x_{N-1,i},u_{N-1,i}) \end{pmatrix}, \quad H(z) := (*_2),$$

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 6.11 shows the exchange of information schematically.

[33] available at www.nmpc-book.com

[34] see www.nonlinearmpc.com

[35] For the sake of comprehensibility we only consider systems in discrete time. In case of continuous

$$\begin{tikzcd}
\boxed{\text{NLP (6.23)}} \arrow[rr, bend left, "{u_{k,i},\, x_0}"] & & \boxed{\text{Dynamics}} \arrow[ll, bend left, "{f(z) \leftarrow x_{k,i}}"]
\end{tikzcd}$$

Figure 6.11: Exchange of information between the optimization problem and the system dynamics

Consequently, in (6.23) after recursive discretization we have

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

$$G(z) := (*_1), \quad H(z) := (*_2),$$

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 $k \in \{0, \ldots, N-1\}$ 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 (6.23) while the C++ software is able to perform

systems, we replace the system dynamics by a numerical approximation, cf. (2.6).

each of the discretization techniques and has links to various optimization packages which can be selected for solving (6.23).

Now that we have defined the optimization problem, we might expect difficulties whenever (6.23) 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 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 they may happen. Often, one can avoid such situations e.g. by adding constraints or chosing 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 4.2 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 2.1 and 3.1), we reconsider the example from this section. In Figure 6.12 we show the optimal open loop trajectories starting in $x_0 = (3, 0.75)^T$ for different N.

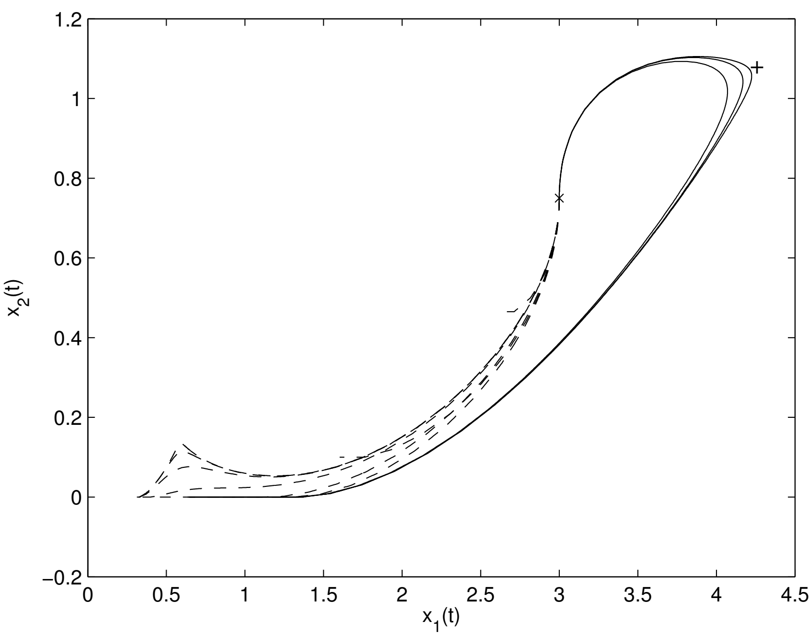

Figure 6.12: Open loop trajectories for the model from Section 4.2 for $x_0 = (3, 0.75)^T$ and various $N \leq 45$ (dashed) and $N \geq 50$ (solid).

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 increase the dimensionality in (6.23) 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 carefully and consider a series of numerical experiments with different parameters in order to verify the validity of their results.

7 Conclusion

In this paper we have demonstrated that NMPC provides an efficient way of numerically solving dynamic decision problems in economics. 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). Our examples show that this method can address deterministic and stochastic model variants with good accuracy as well as models with multiple domains of attraction and thresholds. The NMPC also permits us to compute state and control variables for models where state variables are phased in or phased out. As we have shown, we can also numerically solve models with regime shifts in the dynamics. 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. 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.

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). Thresholds externalities in economic development, *Quarterly Journal of Economics* **105**(2): 501–526.

Barro, R. J., Mankiw, G. and Sala-i Martin, X. (1995). Capital mobility in neoclassical models of growth, *American Economic Review, American Economic Association* **85**(1): 103–115.

Benhabib, J. and Farmer, R. E. A. (1999). Indeterminacy and sunspots in macroeconomics, *in* J. B. Taylor and M. Woodford (eds), *Handbook of Macroeconomics, Volume 1, Part A*, Elsevier, pp. 387–448.

Blanchard, O. J. and Fischer, S. (1989). *Lectures on Macroeconomics*, MIT Press, Cambridge, MA.

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. 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. (2012). An exponential turnpike theorem for dissipative optimal control problems, Preprint, University of Bayreuth.

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.

Eo, Y. (2009). Bayesian analysis of DSGE models with regime switching, Manuscript, Department of Economics, Washington University, St. Louis.

Ernst, E. and Semmler, W. (2012). Labor market, credit market and macroeconomic dynamics, Manuscript, New School for Social Research, New York.

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.

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.

Gallegati, M., Gallegati, M., Ramsey, J. B. and Semmler, W. (2011). The US wage Phillips curve across frequencies and over time, *Oxford Bulletin of Economics and Statistics* **73**: 489–508.

Gilchrist, S. and Zakrajek, E. (2011). Credit spreads and business cycle fluctuations, NBER Working Paper 17021.
URL: *http://www.nber.org/papers/w17021*

Greiner, A., Grüne, L. and Semmler, W. (2012). Economic growth and the transition from non-renewable to renewable energy, SSRN Paper.
URL: *http://dx.doi.org/10.2139/ssrn.2098707*

Greiner, A. and Semmler, W. (2008). *The Global Environment, Natural Resources, and Economic Growth*, Oxford University Press, Oxford, UK.

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* . In Press.
URL: *http://dx.doi.org/10.1016/j.automatica.2012.12.003*

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.

Hartl, R., Kort, P., Feichtinger, G. and Wirl, F. (2000). Multiple equilibria and thresholds due to relative investment costs: non-concave-concave, focus-node, continuous-discontinuous, Working paper, University of Technology, Vienna.

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.

Heal, G. (2009). The economics of renewable energy, NBER Working Paper 15081.
URL: *http://www.nber.org/papers/w15081*

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.

Maurer, H. and Semmler, W. (2010). An optimal control model of oil discovery and extraction., *Appl. Math. Comput.* **217**(3): 1163–1169.

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.

Roch, F. and Uhlig, H. (2012). The dynamics of sovereign debt crises and bailouts, Manuscript, University of Chicago.

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.

Sims, C. A. (2002). Solving linear rational expectations models, *Computational Economics* **20**(1–2): 1–20.

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.