

Asset Pricing with Dynamic Programming*

Lars Grüne

Mathematisches Institut
Fakultät für Mathematik und Physik
Universität Bayreuth
95440 Bayreuth, Germany
`lars.gruene@uni-bayreuth.de`

Willi Semmler

New School, Schwartz Center for Economic Policy Analysis, New York and
Center for Empirical Macroeconomics, Bielefeld University, Germany
`semmlerw@newschool.edu`

November 6, 2006

Abstract: The study of asset price characteristics of stochastic growth models such as the risk-free interest rate, equity premium and the Sharpe-ratio has been limited by the lack of global and accurate methods to solve dynamic optimization models. In this paper a stochastic version of a dynamic programming method with adaptive grid scheme is applied to compute the asset price characteristics of a stochastic growth model. The stochastic growth model is of the type as developed by Brock and Mirman (1972) and Brock (1979, 1982). It has become the baseline model in the stochastic dynamic general equilibrium literature. In a first step, in order to test our procedure, it is applied to this basic stochastic growth model for which the optimal consumption and asset prices can analytically be computed. Since, as shown, our method produces only negligible errors, as compared to the analytical solution, in a second step, we apply it to more elaborate stochastic growth models with adjustment costs and habit formation. In the latter model preferences are not time separable and past consumption acts as a constraint on current consumption. This model gives rise to an additional state variable. We here too apply our stochastic version of a

*A version of this paper has been presented at the Conference on Computation in Economics and Finance, Washington, D.C., June 2005, a Workshop at Chuo University, Tokyo, September 2005, and at a Seminar of the German Bundesbank, June 2006. We want to thank the audience for comments. We are also grateful for extensive comments from a referee.

dynamic programming method with adaptive grid scheme to compute the above mentioned asset price characteristics. We show that our method is very suitable to be used as solution technique for such models with more complicated decision structure.

JEL Classification: C60, C61, C63, D90, G12

Keywords: stochastic growth models, asset pricing, stochastic dynamic programming, adaptive grid

1 Introduction

In this paper we apply stochastic dynamic programming to intertemporal asset pricing models. In the literature we can find several models. The first generation of intertemporal asset pricing models has employed exogenous dividend stream¹ to match financial statistics of the models to the data. Those models have had difficulties to generate adequate characteristics such the risk-free interest rate, equity premium and the Sharpe-ratio, a measure of the risk-return trade-off. In general in those models the risk-free rate turns out to be too high and the mean equity premium and Sharpe-ratio too low compared to what one finds in time series data.

The next generation of asset pricing models has been based on the stochastic growth model with production originating in Brock and Mirman (1972) and Brock (1979, 1982). The Brock approach extends the asset pricing strategy beyond endowment economies to economies that have endogenous state variables including a capital stock that is used in production. Authors, building on this tradition,² have argued that it is crucial how consumption is modelled. In stochastic growth models the randomness occurs in the production function of firms and consumption and dividends are derived endogenously. Yet, the latter type of models have turned out to be even less successful. Given a production shock, consumption can be smoothed through savings and thus asset market facts are even harder to match. Recently numerous extensions of the basic model have been proposed.³

Yet, one of the most important issues in this endeavor to solve asset pricing models has become the accuracy of the solutions of the asset price model. In this paper we argue that we have to be sufficient confident in the accuracy of the solution method before we can apply it to models that are set up to solve the asset pricing puzzles. The conjecture has been that linearization methods to solve the optimal consumption path in feedback form and to use the growth of marginal utility of consumption as discount factor for pricing assets have been insufficient. Other methods have been developed that use second order approximations (Collard and Juillard, 2001), the perturbation method developed by Judd (1996) and the spline method (Trick and Zin, 1997), to name a few. Our attempt here is to discuss and to apply a solution technique, that builds on the the Hamilton-Jacobi-Bellman equation, the dynamic programming method, to address the accuracy issue and to solve asset pricing models.

The stochastic dynamic programming algorithm, with discretization of the state space and adaptive gridding strategy has already appeared useful in other applications.⁴ In

¹Breeden (1979) and Lucas (1978) for example.

²See Rouwenhorst (1995, Akdeniz and Dechert (1997), Jerman (1998), Boldrin, Christiano and Fisher (2001), Lettau and Uhlig (1999) and Hansen and Sargent (2002), the latter in a linear-quadratic economy. The Brock model has also been used to evaluate the effect of corporate income tax on asset prices, see McGrattan and Prescott (2001).

³For a recent account of the gap between models and facts, see Boldrin, Christiano and Fisher (2001), Cochrane (2001, ch. 21), Lettau, Gong and Semmler (2001) and Semmler (2006, chs. 9-10 and 15).

⁴For economic applications of this method, see Grüne and Semmler (2004). They show that solving the stochastic growth models through linearizations for the cases where the value function may have steep curvature, or possibly is non-smooth, may not be appropriate. In Becker, Grüne and Semmler (2006) dynamic programming is compared to the second order approximation of Schmidt-Grohe and Uribe (2004) and it is found that in dynamic programming the errors are independent from the distance to the stationary state whereas for the second order approximation the errors rapidly increases further away

such a method we do not need to use fixed grids, but adaptive space discretization. In some of the literature, see Munos and Moore (2002) and Trick and Ziu (1997), adaptive methods for dynamic programming have been used, but those methods either do not allow the derivation of rigorous error bounds or are computationally very expensive, see the discussion in Sect. 2. In the method applied here efficient and reliable local error estimation is undertaken and used as a basis for a local refinement of the grid in order to deal with the accuracy problem. This procedure allows for a global dynamic solution method of deterministic as well as stochastic intertemporal decision problems.

In order to test the accuracy of such a stochastic dynamic programming algorithm with flexible grid size as solution method for asset pricing models, we numerically provide the global solution and dynamics of a basic stochastic growth model, as based on Brock and Mirman (1972) and Brock (1978, 1982). The basic model can analytical be solved for the policy function in feedback form. Moreover, the asset prices, the risk-free interest rate, the equity premium and the Sharpe-ratio, can, once the model is solved analytically for the sequence of optimal consumption, easily be solved numerically and those solutions compared to the numerical solutions obtained by our stochastic dynamic programming method. As will be shown the errors, as compared to the analytical solutions, are quite small. Our method uses adaptive space discretization. In the method efficient and reliable local error estimation is undertaken and used as a basis for a local refinement of the grid where we are also able to deal with steep slopes or other non-smooth properties of the value function (such as non-differentiability), obtaining quite accurate solutions even further away from some stationary state⁵.

After the method is tested for its accuracy we can then apply it to more complicated extensions of the model which may include different utility functions, in particular habit formation, adjustment cost of capital, a two sector economy, or heterogenous firms and households.⁶ In our extensions we focus here on habit formation and adjustment costs of investment. Our new solution method is applied to such an extended asset pricing model.

Since, as aforementioned, time separable preferences fail to match financial market characteristics an enormous effort has been invested into models with time non-separable preferences, such as habit formation models, which allow for adjacent complementarity in consumption. Past consumption enters here as a constraint, defined as external habit persistence where the aggregate level of consumption serves as a benchmark level, or internal habit persistence where a household's own past consumption is viewed as a benchmark.

There is a long tradition in economic theory where it is assumed that habits are formed through past consumption.⁷ In those models high level of consumption in the past depresses current welfare and high current consumption depresses future welfare. This can

from it.

⁵For the former two problems, see Grüne and Semmler (2004), and for the latter, see Becker, Grüne and Semmler (2006).

⁶For further detailed studies see, for example, Campbell and Cochrane (1999), Jerman (1998), Boldrin, Christiano and Fisher (2001), Cochrane (2001, ch. 21) and Akdezin and Deckert (1997) for extensions along those lines.

⁷Habit persistence is nowadays used to understand a wide range of issues in growth theory (Carrol et al. 1997, 2000, Alvarez-Cuadrado et al. 2004) macroeconomics (Fuhrer, 2002), and business cycle theory (Boldrin et al, 2001). For a first use of habit persistence in a dynamic decision model see Ryder and Heal (1973).

be written as ratios of current over past consumption (Abel 1990, 1999) or in difference form as $(1 - \alpha)c_t + \alpha(c_t - c_{t-1})$ with c_t current, c_{t-1} past consumption and α a respective weight. This form of habit formation will be chosen in this paper.

This type of habit specification gives rise to time non-separable preferences where risk aversion and intertemporal elasticity of substitution are separated and a time variation of risk aversion will arise. If we define surplus consumption as $s_t = \frac{c_t - X_t}{c_t}$ with X_t , the habit, and γ , the risk aversion parameter, then the time variation of risk-aversion is $\frac{\gamma}{s_t}$: the risk aversion falls with rising surplus consumption and the reverse holds for falling surplus consumption. A high volatility of the surplus consumption will lead to a high volatility of the growth of marginal utility and thus to a high volatility of the stochastic discount factor.

Habit persistence in asset pricing has been introduced by Constantinides (1990) in order to study the equity premium problem. Asset pricing models along this line have been further explored by Campbell and Cochrane (1999), Jerman (1998), and Boldrin et al. (2001). Yet, as mentioned before, introducing habit persistence in stochastic models with production may just produce smoother consumption. As the literature has demonstrated (Jerman 1998, and Boldrin et al. 2001) one also needs adjustment costs of investment to obtain better results on the equity premium. By choosing such a model we will not, following Jerman (1998), allow for elastic labor supply, but rather employ a model with fixed labor supply, since the latter, as shown in Lettau and Uhlig (2000), provides the most favorable case for matching the model with the financial market characteristics.

The paper is organized as follows. Section 2 discusses related literature on solution methods. Section 3 presents the stochastic dynamic programming algorithm. Section 4 introduces the stochastic decision problem in asset pricing. Section 5 applies our method to an analytically solvable basic model of asset pricing and studies the accuracy of our solution method. Section 6 introduces habit formation and adjustment costs of investment and solves for the financial characteristics we want to study. Section 7 interprets our results in the context of previous literature on habit formation models. Section 8 concludes the paper.

2 Related Literature on Solution Methods

In the literature on solving asset pricing models one can find a vast amount of different approaches.⁸ We will focus here on those approaches that are employing dynamic programming which is designed to simultaneously approximate both the decision variables and the value function. It is closely related to the Hamilton-Jacobi-Bellman equation. Many of the recent versions of dynamic programming use state-of-the-art mathematical and numerical techniques for making this approach more efficient. Here we apply an adaptive gridding algorithm that works for very general Hamilton-Jacobi-Bellman equations, see Section 3 for details. In the present section we briefly review similar approaches and highlight similarities and differences to our approach.

One of the fundamental difficulties with the dynamic programming approach is that the computational load grows exponentially with the dimension of the problem, a phenomenon

⁸For an early extensive survey of those techniques, see Taylor and Uhlig (1990).

known as the “curse of dimensionality” (see Rust (1996) for a comprehensive account on complexity issues). In our case, for computing asset pricing in the context of stochastic growth models, starting with Brock and Mirman (1972) as suggested in the literature, the problem to be solved is two dimensional, hence this is not a crucial aspect here. Nevertheless, for the sake of completeness we want to mention approaches like randomly distributed grid points (Rust (1997)) or so called low discrepancy grids (Rust (1996), Reiter (1999)) which are able to break the curse of dimensionality. In principle also Monte–Carlo techniques like in Keane and Wolpin (1994) allow for breaking the curse of dimensionality, but as Rust (1997) points out, the specific algorithm in Keane and Wolpin (1994) uses an interpolation technique which again is subject to exponential growth of the numerical cost in the space dimension.

For low dimensional problems the goal of the numerical strategy is not to avoid the curse of dimensionality but rather to reduce the computational cost for a problem of fixed dimension. For this purpose, two main approaches can be found in the literature, namely higher order approximations and adaptive gridding techniques; the latter will be used in our numerical approach.

The idea of high order approximations lies in exploiting the smoothness of the optimal value function: if the optimal value function turns out to be sufficiently smooth, then methods using approximations by smooth functions, like Chebyshev polynomials (Rust (1996), Judd (1996), Jermann (1998)), Splines (Daniel (1976), Johnson et al. (1993), Trick and Zin (1993, 1997)) or piecewise high–order approximations (Falcone and Ferretti (1998)) can be very efficient. Smoothness of the optimal value function is also assumed in the method using second order approximation of the policy function (see Schmidt–Grohe and Uribe 2004 and Collard and Juillard, 2001). Smoothness is also the basis of other high–order strategies, like in finite difference approximations (Candler (2001)), Gaussian Quadrature discretization (Tauchen and Hussey (1991), Burnside (2001)) and in perturbation techniques (Judd (1996)). Yet, the latter should also work if the value function is only piecewise smooth.

Some of these methods (like Spline and piecewise high order approximation) use a (fixed) grid discretization of the state space similar to our approach. The combination of adaptive grids with higher order approximation has been investigated for several one dimensional examples in Bauer, Grüne and Semmler (2006). This combination yields good results when the optimal value function is smooth, however, the resulting numerical scheme may become unstable in the presence of non–differentiable optimal value functions. Since, in addition, the practical implementation of spline interpolation on non–uniform grids in more than one space dimension becomes rather complicated and computationally expensive, we have chosen piecewise linear approximations in this paper.

Concerning discretization techniques it should be noted that from the complexity point of view it turns out to be optimal to solve the dynamic programming problem on successively finer grids, using a one–way multigrid strategy (Chow and Tsitsiklis (1991), see also Rust (1996)). In fact, our adaptive gridding algorithm is similar to this approach since the approximation on the previous grid Γ_i is always used as the initial value for the computation on the next finer adaptive grid Γ_{i+1} . This also explains the large reduction in computation time observed for our approach compared to the computation on *one* fixed equidistant grid.

Let us now turn to our adaptive gridding techniques. Perhaps closest to our approach are the techniques discussed in Munos and Moore (2002). Here a number of heuristic techniques are compared which lead to local and global error indicators which can in turn be used for an adaptive grid generation. Some of the indicators discussed in this paper bear some similarity with our residual based estimator, though rigorous estimates as given below are not given there. In any case, the authors report that these techniques are unsatisfactory and argue for a completely different approach which measures the influence of local errors in certain regions on the global error by analyzing the information flow on the Markov chain related to the discretization of the (deterministic) problem at hand. The reason for this lies in the fact that the model problem treated by Munos and Moore (2002) has a discontinuous optimal value function, which often happens in technical problems with boundary conditions. In fact, also our adaptive scheme performs rather poorly in presence of discontinuities but since our economic problems do always have continuous optimal value functions, Munos' and Moore's conclusions do not apply here. A roughly similar technique is the endogenous oversampling used by Marcat (1994). This is again a heuristic method, which, however, does not lead to adaptive grids but rather selects suitable parts of the state space where the optimally controlled trajectories stay with high probability.

Probably the adaptive approaches with the most solid mathematical background are presented in the papers of Trick and Zin (1993, 1997).⁹ In these papers an alternative approach for the solution of the fully discrete problem is developed using advanced linear programming techniques which are capable of solving huge linear programs with many unknowns and constraints. In Trick and Zin (1993) an adaptive selection of constraints in the linear program is used based on estimating the impact of the missing constraint, a method which is closely related to the chosen solution method but only loosely connected to our adaptive gridding approach. The later paper (Trick and Zin (1997)), however, presents an idea which is very similar to our approach. Due to the structure of their solution they can ensure that the numerical approximation is greater than or equal to the true optimal value function. On the other hand, the induced suboptimal optimal control strategy always produces a value which is lower than the optimal value. Thus, comparing these values for each test point in space one can compute an interval in which the true value must lie, which produces a mathematically concise error estimate that can be used as a refinement criterion. While this approach is certainly a good way to measure errors, which could in particular be less conservative than our measure for an upper bound, we strongly believe that it is less efficient for an adaptive gridding scheme, because (i) the estimated error measured by this procedure is not a local quantity (since it depends on the numerical along the whole suboptimal trajectory), which means that regions may be refined although the real error is large elsewhere, and (ii) compared to our approach it is expensive to evaluate, because for any test point one has to compute the whole suboptimal trajectory, while our residual based error estimate needs only one step of this trajectory.

Let us comment on the idea of a posteriori error estimation. In fact, the idea to evaluate residuals can also be found in the papers of Judd (1996) and Judd and Guu (1997), using,

⁹As mentioned above, this approach also uses splines, i.e., a smooth approximation, but the ideas developed in these papers do also work for linear splines which do not require smoothness of the approximated optimal value function. A comparison of our adaptive approach with splines interpolation can be found in Bauer, Grüne and Semmler (2006).

however, not the dynamic programming operator but the associated Euler equation. In these references the resulting residual is used to estimate the quality of the approximating solution, but to our knowledge it has not been used to control adaptive gridding strategies, and we are not aware of any estimates such as ours which is a crucial property for an efficient and reliable adaptive gridding scheme, particularly needed to solve stochastic problems in asset pricing models.¹⁰

Summarizing our discussion, there are a number of adaptive strategies around which are all reported to show good results, however, they are either heuristic¹¹ and better suited for other classes of problems than asset pricing models or they have nice theoretical features but are practically inconvenient because their implementation is numerically much more expensive than our approach.

3 Stochastic Dynamic Programming

Next we describe the stochastic numerical dynamic programming algorithm that we use to solve the asset pricing characteristics of the stochastic growth model. Details of the adaptive gridding strategy are given in the appendix.

We consider the discrete stochastic dynamic programming equation

$$V(x) = \max_{c \in C} E\{u(x, c, \varepsilon) + \beta(x, \varepsilon)V(\varphi(x, c, \varepsilon))\}. \quad (3.1)$$

Here $x \in \Omega \subset \mathbb{R}^2$, $C \subset \mathbb{R}$, Ω and C are compact sets and ε is a random variable with values in \mathbb{R} . The mappings $\varphi : \Omega \times C \times \mathbb{R} \rightarrow \mathbb{R}^2$ and $g : \Omega \times C \times \mathbb{R} \rightarrow \mathbb{R}$ are supposed to be continuous and Lipschitz continuous in x . Furthermore, we assume that either $\varphi(x, c, z) \in \Omega$ almost surely for all $x \in \Omega$ and all $c \in C$, or that suitable boundary values $V(x)$ for $x \notin \Omega$ are specified, such that the right hand side of (3.1) is well defined for all $x \in \Omega$. The value $\beta(x, \varepsilon)$ is the (possibly state and ε dependent) discount factor which we assume to be Lipschitz and we assume that there exists $\beta_0 \in (0, 1)$ such that $\beta(x, \varepsilon) \in (0, \beta_0)$ holds for all $x \in \Omega$. We can relax this condition if no maximization takes place, in this case it suffices that all trajectories end up in a region where $\beta(x, \varepsilon) \in (0, \beta_0)$ holds. This is the situation for the asset price problem, see the discussion in Cochrane (2001: 27).

Associated to (3.1) we define the dynamic programming operator

$$T : C(\Omega, \mathbb{R}) \rightarrow C(\Omega, \mathbb{R})$$

given by

$$T(W)(x) := \max_{c \in C} E\{u(x, c, \varepsilon) + \beta(x, \varepsilon)W(\varphi(x, c, \varepsilon))\}. \quad (3.2)$$

¹⁰Note that we use combined policy value space iteration. In the approach using second and higher order approximations of the policy functions (see Schmidt-Grohe and Uribe (2004)), the value function has to be approximated in a second step.

¹¹In order to avoid misunderstandings: We do not claim that heuristic methods cannot perform well; in fact they can show very good results. Our main concern about these methods is that one can never be sure about the quality of the final solution of a heuristic method.

The solution V of (3.1) is then the unique fixed point of (3.2), i.e.,

$$T(V) = V. \quad (3.3)$$

For the numerical solution of (3.3) we use a discretization method that goes back to Falcone (1987) in the deterministic case and was applied to stochastic problems in Santos and Vigo–Aguilar (1998). Here we use unstructured rectangular grids: We assume that $\Omega \subset \mathbb{R}^n$ is a rectangular and consider a grid Γ covering Ω with rectangular elements Q_l and nodes x_j and the space of continuous and piecewise multilinear functions

$$\mathcal{W}_\Gamma := \{W \in C(\Omega, \mathbb{R}) \mid W(x + \alpha e_j) \text{ is linear in } \alpha \text{ on each } Q_l \text{ for each } j = 1, \dots, n\}$$

where the e_j , $j = 1, \dots, n$ denote the standard basis vectors of the \mathbb{R}^n , see Grüne (2003) for details of the grid construction. With $\pi_\Gamma : C(\Omega, \mathbb{R}) \rightarrow \mathcal{W}_\Gamma$ we denote the projection of an arbitrary continuous function to \mathcal{W}_Γ , i.e.,

$$\pi_\Gamma(W)(x_j) = W(x_j) \text{ for all nodes } x_j \text{ of the grid } \Gamma.$$

Note that our approach easily carries over to higher order approximations, the use of multilinear approximations is mainly motivated by its ease of implementation, especially for adaptively refined grids.¹² Also, the approach can easily be extended to higher dimensions.

We now define the discrete dynamic programming operator by

$$T_\Gamma : C(\Omega, R) \rightarrow \mathcal{W}_\Gamma, \quad T_\Gamma = \pi_\Gamma \circ T \quad (3.4)$$

with T from (3.2). Then the discrete fixed point equation

$$T_\Gamma(V_\Gamma) = V_\Gamma. \quad (3.5)$$

has a unique solution $V_\Gamma \in \mathcal{W}_\Gamma$ which converges to V if the size of the elements Q_l tends to zero. The convergence is linear if V is Lipschitz on Ω , see Falcone (1987), and quadratic if V is C^2 , see Santos and Vigo–Aguilar (1998).

For the solution of (3.5) as well as for the computation we need to evaluate the operator T_Γ . More precisely, we need to evaluate

$$\max_{c \in C} E\{u(x_j, c, \varepsilon) + \beta(x_j, \varepsilon)W(\varphi(x_j, c, \varepsilon))\}.$$

for all nodes x_j of Γ .

This first includes the numerical evaluation of the expectation E . If ε is a finite random variable then this is straightforward, if ε is a continuous random variable then the corresponding integral

$$\int (u(x, c, \varepsilon) + \beta(x, \varepsilon)V(\varphi(x, c, \varepsilon)))f(\varepsilon)d\varepsilon$$

has to be computed, where f is the probability density of ε . In our implementation we approximated this integral by a trapezoidal rule with 10 equidistant intervals.

¹²The combination of adaptive grids and higher order approximations is currently under investigation.

The second difficulty in the numerical evaluation of T lies in the maximization over c . In our implementation we used Brent's methods with an accuracy of 10^{-3} . Here it turned out that due to the errors induced by the spatial discretization on the grid higher accuracies in Brent's method did not yield higher accuracies in the results.

For the solution of the fixed point equation (3.5) we use the Gauss–Seidel type value space iteration which is described in Section 3 of Grüne (1997) (under the name “increasing coordinate algorithm”), where we subsequently compute $V_{i+1} = S_\Gamma(V_i)$ with S_Γ being a Gauss–Seidel type iteration operator (including the maximization over c) obtained from T_Γ . This iteration is coupled with a policy space iteration: Once a prescribed percentage of the maximizing u -values in the nodes remains constant from one iteration to another we fix all control values and compute the associated value function by solving a linear system of equations using the CGS or BICGSTAB method (in our examples the CGS method turned out to show more reliable convergence behavior). After convergence of this method we continue with the value space iteration using S_Γ until the control values again converge, switch to the linear solver and so on. This combined policy–value space iteration turns out to be much more efficient (often more than 90 percent faster) than the plain Gauss–Seidel value space iteration using S_Γ , which in turn is considerably faster than the Banach iteration $V_{i+1} = T_\Gamma(V_i)$. The gridding algorithm is presented in the appendix.

4 The Stochastic Decision Problem in Asset Pricing

Our stochastic decision problem arising from the stochastic growth model in the Brock tradition which we want to solve and for which we want to compute certain financial measures is as follows. Before we introduce the basic stochastic growth model that we want to apply our solution technique to, see sect. 5, we outline an asset pricing model in a very generic form. The problem we are concerned with is to compute an optimal control, c_t , for the dynamic decision problem

$$V(k, z) = \max_{c_t} E \left(\sum_{t=0}^{\infty} \beta^t u(c_t, X_t) \right) \quad (4.1)$$

with X_t representing other state variables, for example habit, subject to the dynamics

$$\begin{aligned} k_{t+1} &= \varphi_1(k_t, z_t, c_t, \varepsilon_t) \\ z_{t+1} &= \varphi_2(k_t, z_t, c_t, \varepsilon_t) \\ X_{t+1} &= c_t \end{aligned}$$

using the constraints $c_t \geq 0$ and $k_t \geq 0$ and the initial value $k_0 = k, z_0 = z, X_0 = X$. Here $(k_t, z_t, X_t) \in \mathbb{R}^3$ is the state and ε_t are i.i.d. random variables. We abbreviate $x_t = (k_t, z_t, X_t)$ and $\varphi(x, c, \varepsilon) = (\varphi_1(k, z, c, \varepsilon), \varphi_2(k, z, c, \varepsilon), c_t)$, i.e.,

$$x_{t+1} = \varphi(x_t, c_t, \varepsilon_t). \quad (4.2)$$

This optimal decision problem allows for the computation of c in feedback form, i.e. $c_t = c(x_t)$ for some suitable map $c : \mathbb{R}^2 \rightarrow \mathbb{R}$. Based on this c we compute the stochastic

discount factor¹³

$$m(x_t) = \beta \frac{u'(c(x_{t+1}))}{u'(c(x_t))} \quad (4.3)$$

(note that m depends on ε_t and the derivative u' is taken with respect to c_t), which serves as an ingredient for the next step, which consists of solving the asset pricing problem

$$p(x) = E \left(\sum_{t=1}^{\infty} \prod_{s=1}^t m(x_s) d(x_t) \right), \quad (4.4)$$

where $d(x_t)$ denotes the dividend at x_t and $x_0 = x$ and the dynamics are given by

$$x_{t+1} = \varphi(x_t, c(x_t), \varepsilon_t)$$

with c from above.

Finally, we use these values to compute the Sharpe-ratio, which represents the ratio of the equity premium to the standard deviation of the equity return. Hereby R^f is the risk-free interest rate.

$$S = \left| \frac{E(R(x)) - R^f(x)}{\sigma(R(x))} \right| = \frac{-R^f(x) \text{cov}(m(x), R(x))}{\sigma(R(x))}. \quad (4.5)$$

Here

$$R^f(x) = \frac{1}{E(m(x))} \quad (4.6)$$

is the risk-free interest rate and

$$R(x_t) = \frac{d(x_{t+1}) + p(x_{t+1})}{p(x_t)} \quad (4.7)$$

is the gross return.

Note that the equality $E(m(x)R(x)) = 1$ holds, which can serve as a indicator for the accuracy of our numerical solution.

We solve the asset pricing problem in the following three steps: (i) We compute the optimal value function V of the underlying optimal control problem, and compute c from V , (ii) we compute the prices $p(x)$ from c and m , and (iii) we compute the risk-free interest rate, the equity premium and the Sharpe-ratio S from c , m and p .

For our baseline stochastic growth model introduced below, which we solve numerically, both c and p are actually available analytically. This allows us to test each single step of our algorithm by replacing the numerically computed c in (ii) and (iii) and/or p in (iii) by their exact values.

For each of the steps we do now sketch our technique for the numerical computation using the algorithm described above in Section 3.

Step (i):

Compute the approximate optimal value $V_{\Gamma}(x)$ using the method from Section 3 and the adaptive gridding algorithm from the Appendix.

¹³The following financial measures are introduced and studied in detail in Cochrane (2001).

Once V_Γ is computed with sufficient accuracy we obtain the optimal control value $c(x)$ in each point by choosing $c(x)$ such that

$$E(u(c(x)) + \beta V_\Gamma(\varphi(x, c(x), \varepsilon))) = T_\Gamma(V_\Gamma)(x) = \max_c E(u(c) + \beta V_\Gamma(\varphi(x, c, \varepsilon)))$$

holds. Once c is known, m of equ. (4.2) can be computed from this value.

Step (ii):

For computing $p(x)$ we follow the same approach as in Step (i), except that here $c(x)$ is known in advance and hence no maximization needs to be done.

For the computation of p we first solve the dynamic programming equation

$$\tilde{p}(x) = E(d(x) + m(x)p(\varphi(x, c(x), \varepsilon)))$$

which is simply a system of linear equations which we solve using the CGS method. This yields a numerical approximation of the function

$$\tilde{p}(x) = E\left(\sum_{t=0}^{\infty} \prod_{s=1}^t m(x_s) d(x_t)\right)$$

(with the convention $\prod_{s=1}^0 m(x_s) = 1$), from which we obtain p by

$$p(x) = \tilde{p}(x) - d(x).$$

In our numerical computations for the computation of p we have always used the same grid Γ as in the previous computation of V in Step (i). The reason is that it did not seem justified to use a finer grid here, because the accuracy of the entering values c from Step (i) is limited by the resolution of Γ , anyway. However, it might nevertheless be that using a different grid (generated e.g. by an additional adaptation routine) in Step (ii) could increase the numerical accuracy.

Step (iii):

The last step is in principle straightforward, because we do now have all the necessary ingredients to compute the risk-free interest rate, the equity premium and Sharpe-ratio S . However, since all the numerical values entering these computations are subject to numerical errors we have to be concerned with the numerical stability of the respective formulas. The first formula for the Sharpe-ratio, where the numerator represents the equity premium, as the spread between the expected equity return and the risk-free interest rate, is

$$\left| \frac{E(R(x)) - R^f(x)}{\sigma(R(x))} \right|. \tag{4.8}$$

This turns out to be considerably less precise than the second formula

$$\frac{-R^f(x) \text{cov}(m(x), R(x))}{\sigma(R(x))}. \tag{4.9}$$

Since the denominator is the same in both formulas, the difference can only be caused by the different numerators. A further investigation reveals that the numerator of the first formula can be rewritten as

$$R^f(x)(1 - E(m(x))E(R(x)))$$

while that of the second formula reads

$$R^f(x)(E(m(x)R(x)) - E(m(x))E(R(x))).$$

Note that in both formulas we have to subtract values which have approximately the same values, which considerably amplifies the numerical errors. As mentioned above, we know that $E(m(x)R(x)) = 1$, which shows that these formulas are theoretically equivalent. Yet the second formula is more accurate.¹⁴

5 Testing Dynamic Programming for a Basic Model

Next we perform numerical computations for the stochastic growth model as present in Brock and Mirman (1972) and Brock (1979, 1982). We use a basic version as suggested and used by Santos and Vigo-Aguiar (1998) which employs an aggregate capital stock and log utility. The variable X_t in equ. (4.1) will therefore be zero.

The dynamics are defined by

$$\begin{aligned} k_{t+1} &= z_t A k_t^\alpha - c_t \\ \ln z_{t+1} &= \rho \ln z_t + \varepsilon_t \end{aligned}$$

where A, α and ρ are real constants and the ε_t are i.i.d. random variables with zero mean. The return function in (4.1) is $u(c) = \ln c$.

In our numerical computations we used $y_t = \ln z_t$ instead of z_t as the second variable. For our numerical experiments we employed the values

$$A = 5, \quad \alpha = 0.34, \quad \rho = 0.9, \quad \beta = 0.95,$$

¹⁴The higher accuracy of the second formula can be explained as follows: Assume that we have a small systematic additive (deterministic) numerical error in $R(x)$, e.g., $R_{num}(x) \approx R(x) + \delta$. Such errors are likely to be caused by the interpolation process on the grid. Then, using $R^f(x) = 1/E(m(x))$, in the first formula we obtain

$$E(R_{num}(x)) - R^f(x) \approx E(R(x)) - R^f(x) + \delta$$

while in the second formula we obtain

$$\begin{aligned} &R^f(x) \left(E(m(x)R_{num}(x)) - E(m(x))E(R_{num}(x)) \right) \\ &\approx R^f(x) \left(E(m(x)(R(x) + \delta)) - E(m(x))E(R(x) + \delta) \right) \\ &= R^f(x) \left(E(m(x)R(x)) + E(m(x)\delta) - E(m(x))E(R(x)) - E(m(x))\delta \right) \\ &= R^f(x) \left(E(m(x)R(x)) - E(m(x))E(R(x)) \right), \end{aligned}$$

i.e., systematic additive errors cancel out in the second formula.

and ε_t was chosen as a Gaussian distributed random variable with standard deviation $\sigma = 0.008$, which we restricted to the interval $[-0.032, 0.032]$. With these parameters it is easily seen that the computational domain $\Omega = [1, 4] \times [-0.32, 0.32]$ is controlled invariant, i.e., for any $(k, z) \in \Omega$ there exists c such that $\varphi((k, z), c, \varepsilon) \in \Omega$ for all values $\varepsilon \in [-0.032, 0.032]$. Therefore, this Ω can be used as our basis for the state space discretization.

The following figures show some of the numerically obtained results for this example. Figure 5.1 shows the optimal value function, the resulting final adaptive grid with 2977 nodes, and the corresponding asset price p . Figure 5.2 shows the k -components of several optimal trajectories with different initial values. Observe that all trajectories end up near the “equilibrium” point $(k, \ln z) = (2, 0)$ around which they oscillate due to the stochastic influence.




Figure 5.1: Optimal value function V , final adaptive grid and asset price function p




Figure 5.2: k -components of optimal trajectories for different initial values

For our stochastic growth model the optimal control is known analytically and is given by

$$c(x) = c(k, z) = (1 - \alpha\beta)Azk_t^\alpha.$$

Note that c here depends linearly on z .

Since $u'(c) = 1/c$, the stochastic discount factor becomes

$$m(x_t) = \beta \frac{c(x_t)}{c(x_{t+1})}.$$

Furthermore, in this model the dividend is given by $d(x) = c(x)$, and from this one easily computes¹⁵

$$p(x) = \frac{\beta}{1 - \beta} c(x) = \frac{\beta}{1 - \beta} (1 - \alpha\beta) A z k_t^\alpha.$$

Using these values, we can compute the Sharpe-ratio by evaluating the respective expressions numerically. We obtain the value

$$S = 0.007999.$$

Note that the value is independent of x , but the component parts defining S are not constant at all. In Figure 5.3 the values $R^f(x) = 1/E(m(x))$, $\sigma(m(x))$ and $-R^f(x)\text{cov}(m(x), R(x))$, defining the numerator of the Sharpe-ratio, are shown depending on x , all computed from the exact values.




Figure 5.3: Functions $R^f(x)$, $\sigma(m(x))$ and $-R^f(x)\text{cov}(m(x), R(x))$

In addition, Figure 5.4 shows a larger plot of the risk free interest rate $R^f(x)$ together with a plane at level $R^f = 1$, which allows to identify where the rate is below and where above one. Note that in the equilibrium $x = (k, \ln z) = (2, 0)$, around which the optimal trajectories end up in the long run, the value of R^f is 1.06.

¹⁵see also Brock (1982).




Figure 5.4: The risk free interest rate $R^f(x)$ and a plane at level $R^f = 1$

For our numerical tests we used three different accuracies. For a test with lower accuracy we used a discretization of the set C of control values with discretization step size $6.25 \cdot 10^{-2}$ and an adaptive grid with 2161 nodes. These computations were also used for the Figures 5.1 and 5.2 of the numerical examples (visually the results with higher accuracy look exactly the same).

For the test with medium accuracy we used an adaptive grid with 15020 nodes and for the fine accuracy we used 109533 nodes. Table 5.1 shows the respective results with the errors measured in the L_∞ norm.

| accuracy | error in V | error in c | error in p | error in S |
|----------|---------------------|---------------------|---------------------|---------------------|
| coarse | $3.4 \cdot 10^{-3}$ | $7.4 \cdot 10^{-2}$ | 1.5 | $4.8 \cdot 10^{-3}$ |
| medium | $7.3 \cdot 10^{-4}$ | $3.5 \cdot 10^{-2}$ | $6.5 \cdot 10^{-1}$ | $2.9 \cdot 10^{-3}$ |
| fine | $1.6 \cdot 10^{-4}$ | $9.6 \cdot 10^{-3}$ | $1.9 \cdot 10^{-1}$ | $7.9 \cdot 10^{-4}$ |

Table 5.1: Approximation errors for the fully numerical approach

As an alternative, we used the same grids but using the exact $c(x)$ in the grid points with interpolation in between for all our computations (i.e., for p and S). This rules out the errors caused by the maximization routine and leaves us with the pure interpolation errors on the grid. The corresponding accuracies are summarized in Table 5.2.

| accuracy | error in V | error in c | error in p | error in S |
|----------|--------------|---------------------|---------------------|---------------------|
| coarse | – | $2.2 \cdot 10^{-3}$ | $4.2 \cdot 10^{-2}$ | $1.8 \cdot 10^{-4}$ |
| medium | – | $4.8 \cdot 10^{-4}$ | $1.1 \cdot 10^{-2}$ | $4.9 \cdot 10^{-5}$ |

Table 5.2: Approximation errors using exact c in the grid points

From the tables it is interesting to see how the original numerical errors in V , c and p

propagate through the consecutive numerical approximations. In particular, it turns out that the error in S is much smaller than the error in p and about one order of magnitude smaller than the error in c ; in fact in all examples we investigated it is very similar to the error in the underlying optimal value function V from which c is obtained.

We also see that for the exact c even on the coarser grid very good approximations for S are obtained. From our experiments it seems that the crucial point in the numerical procedure is the numerical accuracy for the underlying V and c , while the accuracy of p does not seem to be crucial. Note that the error in the optimal consumption affects the asset price computation in two ways: on the one hand errors in c affect the stochastic discount factor and on the other hand errors in c cause errors in the optimal paths. This combination of two effects seems to be the main cause for the rather large numerical errors in p . On the other hand, the interpolation in the computation of p does not seem to be crucial for the accuracy, as the results in Table 5.2 show. Summarizing, if the consumption strategy c is either known exactly or if it can be determined numerically with a sufficiently high accuracy, then our numerical approach can be expected to be successful.

Note that the errors given above are obtained by maximizing the errors on a 250x250 grid covering the whole computational domain Ω . In a vicinity of the equilibrium the errors are smaller; for instance, the errors in the Sharpe-ratio S are $1.4 \cdot 10^{-3}$, $9.3 \cdot 10^{-4}$ and $3.7 \cdot 10^{-4}$ for the fully numerical approach with coarse, medium and fine grid, respectively.¹⁶ Furthermore, we discuss absolute errors here. The reason for this lies in the fact that in our examples the stochastic process determining the magnitude of the Sharpe-ratio does not seem to affect the accuracy, as we will demonstrate in our next numerical test.

We should mention that in this examples for $\sigma = 0.008$, for the fully numerical approach the errors for the computation of the Sharpe-ratio S for coarse and medium grids are almost of the same order as the Sharpe-ratio itself from which one might conclude that they do not provide a reasonable approximation. In this case, the relative error for S would be of the order of one, i.e. the numerical results could be considered useless. Yet, we would like to explore whether this undesirable behavior remains true if we modify our model parameters in such a way that the Sharpe-ratio (which is unrealistically small as compared to the empirical Sharpe-ratio¹⁷) increases.

To this end we changed our random variables ε_i to a Gaussian distributed random variable with standard deviation¹⁸ $\sigma = 0.018$, which we now restricted to the interval $[-0.072, 0.072]$. In order not to increase the computational domain Ω we reduced the parameter ρ to $\rho = 0.5$ which ensures that $\Omega = [1, 4] \times [-0.32, 0.32]$ remains controlled invariant.

The exact value of S for these parameters evaluates to

$$S = 0.017994.$$

¹⁶We want to note here, that this approximation does not rise much neither for the Sharpe-ratio, when we depart for the equilibrium. This is not so in methods that use the second (or higher) order approximations, as for example employed in Schmidt-Grohe and Uribe (2004), where the error rapidly rises with the distance to the equilibrium, for details of a comparison of the solution from that methods and a dynamic programming solution, see Becker, Grüne and Semmler (2006).

¹⁷For a report on empirically realistic Sharpe-ratios, see for example, see Section 7.

¹⁸The standard deviation of $\sigma = 0.018$ is suggested in the RBC literature, see Boldrin, Christiano and Fisher (2001).

For this set of parameters we have repeated the fully numerical computation on an adaptive grid with 2529 nodes, which is roughly the same amount of nodes as for the coarse accuracy, above. Table 5.3 shows the respective numerical accuracies for $\sigma = 0.018$ compared to the accuracies for the model with lower standard deviation $\sigma = 0.008$ from Table 5.1.

| σ | error in V | error in c | error in p | error in S |
|----------|---------------------|---------------------|---------------------|---------------------|
| 0.008 | $3.4 \cdot 10^{-3}$ | $7.4 \cdot 10^{-2}$ | 1.5 | $4.8 \cdot 10^{-3}$ |
| 0.018 | $2.4 \cdot 10^{-3}$ | $4.0 \cdot 10^{-2}$ | $7.9 \cdot 10^{-1}$ | $3.2 \cdot 10^{-3}$ |

Table 5.3: Approximation for fully numerical approach and different σ

It turns out that the absolute numerical error in the approximation S does not increase when σ increases, implying that the numerical accuracy of our procedure appears to be independent of the standard deviation σ of the underlying random process and thus of the magnitude of the Sharpe-ratio. In other words, when S increases due to larger shocks, then the relative error in S can be expected to decrease.¹⁹

Summing up, our intention so far was to introduce and apply a stochastic dynamic programming algorithm to a basic asset pricing model in order to provide global and reliable solutions. In order to test our algorithm we have applied it to a basic stochastic growth model for which consumption, asset prices and the Sharpe-ratio can analytically be computed. Overall, as we have shown, the errors depend on the discretization step size and grid refinement but are less impacted by the variance of the stochastic component of the model. Our computations show that the optimal consumption, the value function and the Sharpe-ratio can be computed with small absolute errors. We, in particular, could show that the error in the computation of the Sharpe-ratio is independent of the standard deviation of the underlying random process. Our method remains accurate for large shocks and the relative error decreases with increasing Sharpe-ratio.²⁰ Overall the results presented here are very encouraging and our method can be applied to extended versions of the stochastic growth model with some confidence.

6 Solving the Habit Formation Model with Dynamic Programming

Next, we use our numerical scheme to solve the habit formation model with adjustment cost of investment. The model is given by

¹⁹Other recent literature has also evaluated the effect of an increase in σ on the computational accuracy of approximation methods. Collard and Juillard (2001) find large errors if one applies first order approximations to the policy function, yet less so for second order approximations. Second order approximations, however, significantly lose accuracy with larger shocks, see Becker, Grüne and Semmler (2006).

²⁰Note that this result is important since in empirical research a Sharpe-ratio is obtained which is by a factor of 40 to 50 greater than our smallest Sharpe-ratio, the Sharpe-ratio resulting from $\sigma = 0.008$, see Boldrin, Christiano and Fisher (2001) and section 7 of this paper.

$$k_{t+1} = \varphi_1(k_t, z_t, c_t, \varepsilon_t) = k_t + \frac{k_t}{1-\varphi} \left[\left(\frac{I_t}{k_t} \right)^{1-\varphi} - 1 \right]$$

$$\ln z_t = \varphi_2(k_t, z_t, c_t, \varepsilon_t) = \rho \ln z_t + \varepsilon_t,$$

where $I_t = z_t A k_t^\alpha - c_t$, where in our numerical computations we used the variable $y_t = \ln z_t$ instead of z_t as the second variable.

The utility function in (4.1) is now given by

$$u(c_t, X_t) = \frac{(c_t - bX_t)^{1-\gamma} - 1}{1-\gamma}$$

for $\gamma \neq 1$ and by

$$u(c_t, X_t) = \ln(c_t - bX_t)$$

for $\gamma = 1$. Since we are working with internal habit, in our case, we have $X_t = C_{t-1}$.

For our numerical experiments we employed the same parameter values as in section 5. Here too ε_t was chosen as a Gaussian distributed random variable with standard deviation $\sigma = 0.008$, which we restricted to the interval $[-0.032, 0.032]$. With this choice of parameters it is easily seen that the interval $[-0.32, 0.32]$ is invariant for the second variable y_t .

Motivated by our study in Sect. 5 we would like to solve our problem for k_t on an interval of the form $[a, b]$; here we used $[0.1, 10]$. However, the habit persistence implies that for a given habit X_t only those value c_t are admissible for which $c_t - bX_t > 0$ holds, which defines a constraint from below on c_t depending on the habit X_t . On the other hand, the condition $I_t \geq 0$ defines a constraint from above on c_t depending on k_t and $y_t = \ln z_t$. As a consequence, there exist states $x_t = (k_t, y_t, X_t)$ for which the set of admissible control values c_t is empty, i.e., for which the problem is not feasible. On the one hand, we want to exclude these points from our computation, on the other hand we want to have a computational domain which is of a simple shape. A solution for this problem is given by a coordinate transformation which transforms a suitable set $\tilde{\Omega}$ of feasible points to the set $\Omega = [0.1, 10] \times [-0.32, 0.32] \times [0, 7]$ on which we perform our computation. The coordinate transformation $\Psi : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ we use for this purpose shifts the k -component of a point $x = (k, y, X)$ in such a way that the non-feasible points are mapped to a point $\tilde{x}_t \notin \tilde{\Omega}$. It is given by

$$\Psi(k, y, X) := (k - s(y, X), y, X)$$

with

$$s(y, X) = \left(\frac{s_0 + bX}{\exp(y)A} \right)^{\frac{1}{\alpha}} - 0.1$$

where $s_0 = 0.1^\alpha \exp(-0.32)A$ is chosen such that for $y = -0.32$ and $X = 0$ the coordinate change is the identity. This map is built in such a way that for all points $x_t \in \tilde{\Omega} = \Psi^{-1}(\Omega)$ a value c_t with $c_t - bX_t \geq s_0$ is admissible. Note that forward invariance of $\tilde{\Omega}$ is not automatically guaranteed by this construction and indeed there are parameter combinations for which this property does not hold.

This coordinate transformation allows us to set up our dynamic programming algorithm on a set of feasible points without having to deal with a complicated domain $\tilde{\Omega}$, because numerically we can now work on the simple set Ω using the transformed dynamics

$$x_{t+1} = \Psi \circ \varphi(\Psi^{-1}(x_t), c_t, \varepsilon_t)$$

instead of (4.2).

In addition to the parameters specified above, for our numerical experiments we have used the following sets of parameters:

- (a) $\varphi = 0, \quad \gamma = 1, \quad b = 0$
- (b) $\varphi = 0, \quad \gamma = 1, \quad b = 0.5$
- (c) $\varphi = 0, \quad \gamma = 3, \quad b = 0.5$
- (d) $\varphi = 0.8, \quad \gamma = 1, \quad b = 0.5$
- (e) $\varphi = 0.8, \quad \gamma = 3, \quad b = 0.5$

Note that (a) corresponds to the setting of section 5.

Our first set of numerical results shows the behavior of the k_t -component of the optimal trajectories of the optimal control problem (4.1) in Figure 6.1 (a)–(e), the stochastic discount factor along these trajectories in Figure 6.2(a)–(e) and the consumption in Figure 6.3(a)–(e). For all trajectories we used the initial value $(k_0, y_0, X_0) = (2, 0, 4)$ which is near the point around which the optimal trajectory oscillates. Furthermore, for all trajectories we have used the same sequence of the random variables ε_t .

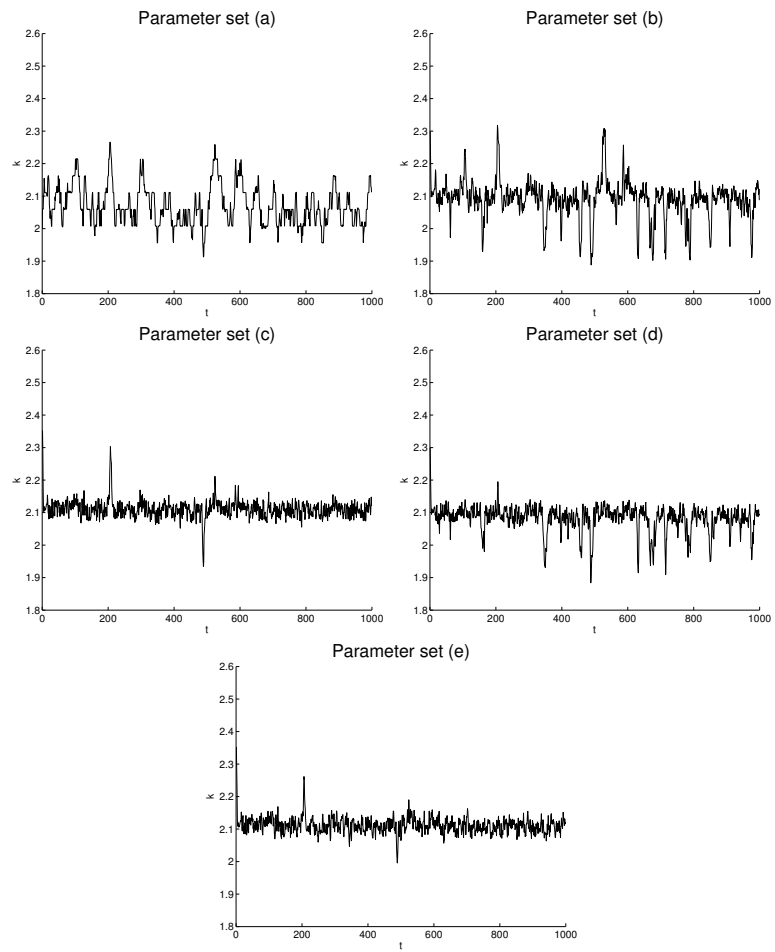

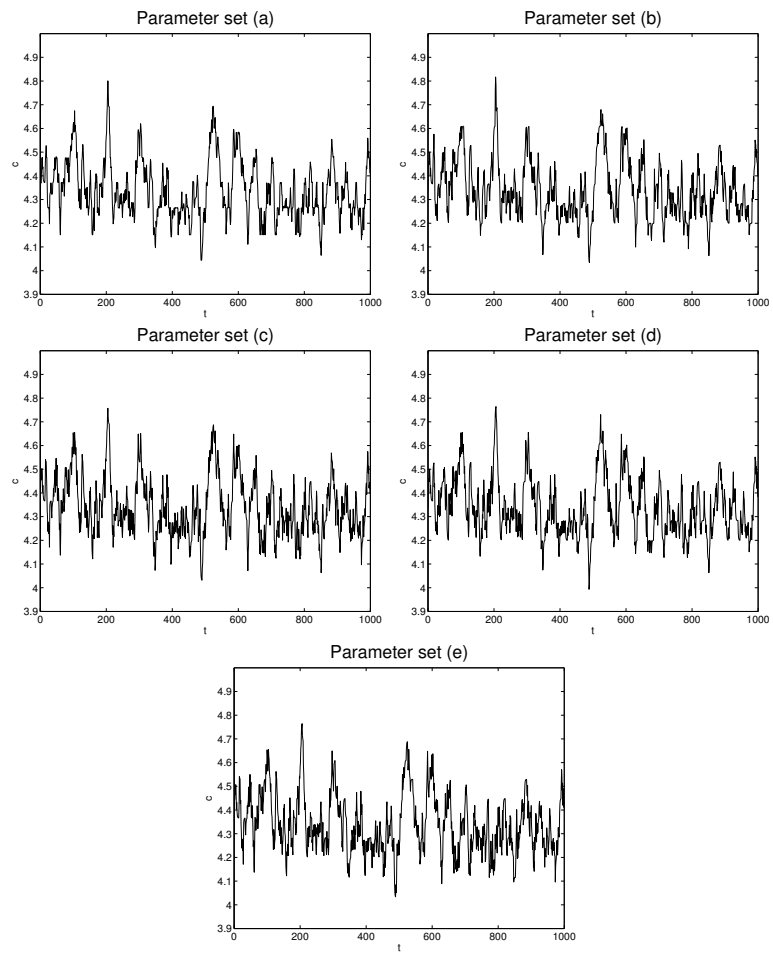


Figure 6.1: k -component of optimal trajectories

Figure 6.2: Stochastic discount factor β along optimal trajectories

Figure 6.3: Consumption c along optimal trajectories

The following Table 6.1 shows several characteristic values obtained from our algorithm. The values were obtained by averaging the values along the optimal trajectories described above.

| parameter set | Sharpe-ratio | equity premium | risk free interest rate |
|---------------|--------------|----------------|-------------------------|
| (a) | 0.0085 | 0.00008 | 1.051 |
| (b) | 0.0153 | 0.00021 | 1.049 |
| (c) | 0.0540 | 0.00227 | 1.084 |
| (d) | 0.0201 | 0.00031 | 1.060 |
| (e) | 0.0572 | 0.00329 | 1.085 |

Table 6.1: Numerically computed values

As can be observed from the figures 6.2 (b) and (d) the volatility of the stochastic discount factor is increased due to habit formation and time varying risk aversion. It is most volatile for the combination of a high γ and habit persistence, see figure 6.2 (c) and (e). In the latter two cases the curvature of the welfare function affects, as one would expect, the volatility of the discount factor. Yet, as the figures 6.3 (a)-(b) show, the consumption path itself is only very little affected by habit persistence and adjustment costs of capital.²¹ From the table 6.1 we can observe that the Sharpe-ratio and equity premium increase strongly with habit persistence and adjustment costs, though not sufficiently to match empirical facts.²² The risk free interest rate is still too high. In particular, as is also well known from asset pricing studies with power utility, the risk free rate rises with higher parameter of relative risk aversion, γ , see table 6.1 (c) and (e).

7 Evaluating the Results on Habit Formation

It is interesting to compare our numerical results that we have obtained, to previous quantitative studies undertaken for habit formation. We in particular will restrict ourselves to a comparison with the results obtained by Boldrin et al. (2001) and Jerman (1998).

Whereas Boldrin et al. use a model with log utility for internal habit, but endogenous labor supply in the household's preferences, Jerman studies the asset price implication of a stochastic growth model, also with internal habit formation but, as in our model, labor effort is not a choice variable. All three papers Boldrin et al. (2001), Jerman (1998) and our variant use adjustment costs of investment in the model with habit formation. Both previous studies claim that habit formation models with adjustment costs come closer to match the financial characteristics of the data.

Yet, there are two reasons why the results of those studies differ from ours. First, in comparison to their parameter choice we have chosen parameters that have often been

²¹Note that our result on habit persistence is a result that Lettau and Uhlig (2000) have also predicted.

²²This result presumably comes from the fact that consumption volatility does not rise sufficiently and thus does not sufficiently covary with the asset returns (which is needed to have a higher equity premium and Sharpe-ratio).

used for the study of the basic stochastic growth models²³ and appear to us less extreme. Table 7.1 reports the parameters and the results.

Both, the study by Jerman (1998) and Boldrin et al. (2001) have chosen a parameter, $\varphi = 4.05$, in the adjustment costs of investment, a very high value which is at the very upper bound found in the data.²⁴ Since the parameter φ smoothes the fluctuation of the capital stock and makes the supply of capital very inelastic, we have rather worked with a $\varphi = 0.8$ in order to avoid such strong volatility of returns generated by high φ . Moreover, both papers use a higher parameter for past consumption, b , than we have chosen. Both papers have also selected a higher standard deviation of the technology shock. Boldrin et al. take $\sigma = 0.018$, and Jerman takes a $\sigma = 0.01$, whereas we use $\sigma = 0.008$ which has been employed in many models.²⁵ Those parameters increase the volatility of the stochastic discount factor, a crucial ingredient to raise the equity premium and the Sharpe-ratio.

| Boldrin et al. ^{a)} | Jerman ^{b)} | Grüne and Semmler ^{c)} | US Data ^{d)} (1954-1990) |
|--|---|---|--------------------------------------|
| $b= 0.73-0.9$ $\varphi=4.15$ $\sigma= 0.018$ $\rho= 0.9$ $\beta= 0.999$ $\gamma= 1$ | $b= 0.83$ $\varphi=4.05$ $\sigma= 0.01$ $\rho= 0.99$ $\beta= 0.99$ $\gamma= 5$ | $b= 0.5$ $\varphi= 0.8$ $\sigma= 0.008$ $\rho= 0.9$ $\beta= 0.95$ $\gamma= 1(3)$ | |
| $R^f = 1.2$ | $R^f = 0.81$ | $R^f = 24.0(34.0)$ | $R^f = 0.8$ |
| $E(R) - R^f = 6.63$ | $E(R) - R^f = 6.2$ | $E(R) - R^f = 1.32$ | $E(R) - R^f = 6.18$ |
| $SR= 0.36$ | $SR= 0.33$ | $SR=0.11$ | $SR=0.35$ |

- a) Boldrin et al.(2001) use a model with endogenous labor supply, log utility for habit formation and adjustment costs, σ quarterly, return data and Sharpe-ratio are, in percentage terms annualized.
- b) Jerman (1998) uses a model with exogenous labor supply, habit formation with coefficient of RRA of 5, and adjustment costs, σ quarterly, return data and Sharpe-ratio are annualized.
- c) This represents the parameters of the cases (d) and (e). The return data and the Sharpe-ratio are annualized.
- d) The following financial characteristics of the data are reported in Jerman (1998), annualized returns and Sharpe-ratio in percentage terms.

Table 7.1: Habit formation models

Jerman, in addition, takes a very high parameter of relative risk aversion, a $\gamma = 5$, which also increases the curvature of the welfare function, the volatility of the discount factor and the equity premium when used for the pricing of assets. Jerman also takes a much higher persistence parameter for the technology shocks, a $\rho = 0.99$, from which one knows that

²³See Santos and Vigo-Aguiar (1998).

²⁴See for example, Kim (2002) for a summary of the empirical results reported on φ in empirical studies.

²⁵This value of σ has also been used by Santos and Vigo-Aguiar (1998).

it will make the stochastic discount factor more volatile too. All in all, both studies have chosen parameters which seem to bias the results toward the empirically found financial characteristics.

Second, the studies by Jerman (1998) and Boldrin et al (2001) have employed different procedures to solve the intertemporal decision problem. Boldrin et al. use the Lagrangian multiplier from the corresponding planner's problem to solve for asset prices. The procedure is based on Christiano and Fisher (2000). Jerman uses a log-linear approach to solve the model and the procedure is compared to non-linear approximation. Concerning the accuracy of different solution procedures there are intricate issues involved such as approximation errors due to the order of approximation, the role of parameter choice (steeper curvature creates greater errors), the distance to the equilibrium, absolute or relative errors and so on.²⁶ We also want to note that there is a crucial constraint in habit formation models, namely that the surplus consumption has to remain non-negative when the optimal solution, c_t , is computed.²⁷ As we have shown in Section 6 this constraint has to be treated properly in the numerical solution method.

Moreover, we want to note, since the risk aversion, $\frac{\gamma}{s_t}$ for power utility, falls with the rising consumption surplus ratio, given by $s_t = \frac{c_t - X_t}{c_t}$, the habit persistence model predicts a rising risk aversion (rising Sharpe-ratio) in recessions and falling risk aversion (falling Sharpe-ratio) in booms, for details see Cochrane (2001: 471). Thus, the risk aversion and the Sharpe-ratio move over the business cycle. A state dependent risk aversion and Sharpe-ratio were also observable in our computations. What we have depicted in table 7.1 are averages along the optimal trajectories in the neighborhood of the steady state.

A further remark is needed on the high risk free rate. In table 6.1 it is computed as $1 + R^f$ and for quarterly data. In table 7.1, we have R^f for annualized data, in percent terms. As table 7.1 shows the risk free rate is highest for the high γ , and with this the annual risk free rate can go up to 34 percent. (In the basic model with $\gamma = 1$ and no habit persistence and to adjustment cost, R^f is about 24 percent.) This is much too high as compared to empirical data where the annual risk free rate is usually reported to be about 1 to 2 percent (depending on time periods considered). Yet, as well known in the literature²⁸ the equity premium and Sharpe-ratio are favorably impacted by a higher γ which, however, produces also a higher risk free rate. Thus, an attempt to increase the equity premium and Sharpe-ratio in the context of the above presumed utility function will make the risk free rate moving in the wrong direction, namely it rises too. We also want to note that the high risk free rate is computed by using a low β . We have chosen $\beta = 0.95$, as compared to a $\beta = 0.99$ and $\beta = 0.999$ in the study by Boldrin et al. and Jerman respectively. As also well known in the literature a higher β will reduce again the risk free rate.²⁹

Overall, one is inclined to conclude that previous studies, because of the specific parameter

²⁶For further discussion on those issues, see also Collard and Juillard (2001), in particular on the order of approximation and the role of parameters.

²⁷Boldrin et al. (2001:154) just make a general statement "that $C_t \leq bC_{t-1}$...[is] never observed in the Monte Carlo simulations..."

²⁸See Cochrane (2001, ch. 21) and Campbell, Lo and MacKinley (1997, ch. 8.2).

²⁹We have taken a low value of β in order to make our study comparable to previous studies solving the stochastic growth model with dynamic programming. A too high risk free interest rate can always be reduced by a higher β , see Cochrane (2001, ch. 21) and Campbell, Lo and MacKinley (1997, ch. 82.).

choice, have obtained more favorable results in solving the equity premium puzzle. Though also our results for the habit formation model point in the same direction, we show by using a method that has been tested for accuracy, that there are still puzzles remaining for the consumption-based asset pricing models. Finally, we want to note that in our study we have chosen a model variant with no endogenous labor supply, which, as Lettau and Uhlig (2000) show, is the most favorable model for asset pricing in a production economy, since including labor supply as a choice variable, would even reduce the equity premium and the Sharpe-ratio.

8 Conclusions

In recent times extensive research effort has been devoted to study the asset price characteristics, such as the risk-free interest rate, the equity premium and the Sharpe-ratio, based on the stochastic growth model. Although it includes production this is a consumption-based asset pricing model. The failure of the model with simple preferences, such as power utility, to match the empirical characteristics of asset prices and returns has given rise to the exploration of numerous extensions of the basic model that allows for different preferences and technology shocks, adjustment cost of capital, two sector economies and heterogenous firms³⁰ and households.³¹

The aim of this paper was two-fold. First, we wanted to apply a solution procedure to solving asset pricing models that has passed an accuracy test. For this purpose we used a basic stochastic growth model and apply our dynamic programming algorithm to that basic model for which consumption, asset prices and the Sharpe-ratio can analytically be computed. We show, the errors of the equity premium and Sharpe-ratio depend on the discretization step size and grid refinement but not on the variance of the shocks. Our computations show that the optimal consumption, the value function and the Sharpe-ratio can be computed with small absolute errors. We also show that the relative errors in computing the Sharpe-ratio can be small. The error in the computation of the Sharpe-ratio is independent of the magnitude of the Sharpe-ratio. In other words the relative error decreases with increasing Sharpe-ratio.³² We also want to note that our method appears to more accurate than the equidistant spline method and the approximation methods using second order approximations of the policy function.³³

Next we study the asset pricing characteristics of a model with more promising extensions than the basic growth model by using our numerical method that appears to solve the asset pricing models with higher accuracy. As extended model we have chosen a model with more complex decision structure, a model with habit persistence, and augmented it, along the line of Boldrin et al (2001) and Jerman (1998), with adjustment costs of

³⁰A model with heterogenous firms in the context of a Brock type stochastic growth model can be found in Akdeniz and Dechert (1997) who are able to match, to some extent, the equity premium by building on idiosyncratic stochastic shocks to firms.

³¹For a model with heterogenous households, see Constantinides and Duffie (1996).

³²Note that in empirical research a Sharpe-ratio is obtained which is by a factor of 40 to 50 greater than our smallest Sharpe-ratio, the Sharpe-ratio resulting from $\sigma = 0.008$, see Boldrin, Christiano and Fisher (2001).

³³For a comparison of our method with the former, see Grüne, Bauer and Semmler (2006) and for the latter, see Becker, Grüne and Semmler (2006).

investment. Extensive research effort has recently been devoted to such a model variant in order to match the empirical characteristics of asset prices and returns. Our results, based on our dynamic programming algorithm shows that, even if habit persistence is jointly used with adjustment costs of investment, there are still puzzles remaining for consumption based asset pricing models.³⁴

³⁴A new promising direction is to include loss aversion in preferences, see Grüne and Semmler (2005), which gives a much better result on the equity premium and Sharpe-ratio.

9 Appendix: Details of the Gridding Strategy

The basic idea of our adaptive gridding algorithm lies in evaluating the residual of the operator T applied to V_Γ , as made precise in the following definition. Here for any subset $B \subset \Omega$ and any function $W \in C(\Omega, \mathbb{R})$ we use

$$\|W\|_{\infty, B} := \max_{x \in B} |W|.$$

(i) We define the *a posteriori error estimate* η as a continuous function $\eta \in C(\Omega, \mathbb{R})$ by

$$\eta(x) := |T(V_\Gamma)(x) - V_\Gamma(x)|.$$

(ii) For any element Q_l of the grid Γ we define the *elementwise error estimate*

$$\eta_l := \|\eta\|_{\infty, Q_l}$$

(iii) We define the *global error estimate* η_{\max} by

$$\eta_{\max} := \max_l \eta_l = \|\eta\|_{\infty}.$$

It is shown in Grüne (2003), that for this error estimate the inequalities

$$\frac{\eta_{\max}}{1 + \beta_0} \leq \|V - V_\Gamma\|_{\infty} \leq \frac{\eta_{\max}}{1 - \beta_0}$$

holds. These inequalities show that the error estimate is *reliable* and *efficient* in the sense of numerical error estimator theory, which is extensively used in the numerical solution of partial differential equations. Furthermore, $\eta(x)$ is continuous and one can show that a similar upper bound holds for the error in the derivative of V and V_Γ .

If the size of a grid element tends to zero then also the corresponding error estimate tends to zero, even quadratically in the element size if V_Γ satisfies a suitable “discrete C^2 ” condition, i.e., a boundedness condition on the second difference quotient.

This observation shows that refining elements carrying large error estimates is a strategy that will eventually reduce the element error and consequently the global error, and thus forms the basis of the adaptive grid generation method which we will describe in the next section.

Clearly, in general the values $\eta_l = \max_{x \in Q_l} \eta(x)$ can not be evaluated exactly since the maximization has to be performed over infinitely many points $x \in Q_l$. Instead, we approximate η_l by

$$\tilde{\eta}_l = \max_{x_T \in X_T(Q_l)} \eta(x_T),$$

where $X_T(Q_l)$ is a set of test points. In our numerical experiments we have used the test points indicated in Figure 9.1.




Figure 9.1: Test points $X_T(Q_l)$ for a 2d element Q_l

The adaptive grid itself was implemented on a tree data structure in the programming language C. The adaptive refinement follows the standard practice in numerical schemes and works as follows:

- (0) Choose an initial grid Γ_0 , set $i = 0$, fix a refinement threshold $\theta \in (0, 1)$
- (1) Compute V_{Γ_i} and the (approximated) error estimates $\tilde{\eta}_l$ and $\tilde{\eta}_{\max}$. If a desired accuracy or a maximally allowed number of nodes is reached, then stop
- (2) Refine all elements Q_l with $\tilde{\eta}_l \geq \theta \tilde{\eta}_{\max}$, denote the new grid by Γ_{i+1}
- (3) Set $i := i + 1$ and go to (1)

Here for the solution of V_{Γ_i} for $i \geq 1$ we use the previous solution $V_{\Gamma_{i-1}}$ as the initial value for the iteration described in Section 3, which turns out to be very efficient.

During the adaptation routine it might happen that the error estimate causes refinements in regions which later turn out to be very regular. It is therefore advisable to include a coarsening mechanism in the above iteration. This mechanism can, e.g., be controlled by comparing the approximation V_{Γ_i} with its projection $\pi_{\tilde{\Gamma}_i} V_{\Gamma_i}$ onto the grid $\tilde{\Gamma}_i$ which is obtained from Γ_i by coarsening each element once. Using a specified coarsening tolerance $tol \geq 0$ one can add the following step after Step (2).

- (2a) Coarsen all elements Q_l with $\tilde{\eta}_l < \theta \tilde{\eta}_{\max}$ and $\|V_{\Gamma_i} - \pi_{\tilde{\Gamma}_i} V_{\Gamma_i}\|_{\infty, Q_l} \leq tol$.

This procedure also allows to start from rather fine initial grids Γ_0 , which have the advantage of yielding a good approximation $\tilde{\eta}_l$ of η_l . Unnecessarily fine elements in the initial grids will this way be coarsened afterwards.

In addition, it might be desirable to add additional refinements in order to avoid large differences in size between adjacent elements, e.g., to avoid degeneracies. Such regularization steps could be included as a step (2b) after the error based refinement and coarsening has been performed. In our implementation such a criterion was used; there the difference in refinement levels between two adjacent elements was restricted to at most one. Note that the values in the hanging nodes (these are the nodes appearing at the interface between two elements of different refinement level) have to be determined by interpolation in order to ensure continuity of V_{Γ} .

In addition, our algorithm allows for the anisotropic refinement of elements: consider an element Q of Γ (we drop the indices for notational convenience) and let $X_{new,i}$ be the set of potential new nodes which would be added to Γ if the element Q_l was refined in coordinate direction e_i , see Figure 9.2.




Figure 9.2: Potential new nodes $X_{new,1}$ (left) and $X_{new,2}$ (right) for a 2d element Q

Define the error estimate in these nodes for each coordinate direction e_i by $\eta_{dir,i} := \max_{x \in X_{new,i}} \eta(x)$ and define the overall error measured in these potential new nodes by $\eta_{dir} := \max_{i=1,\dots,n} \eta_{dir,i}$. Note that $\eta_{dir} \leq \eta_l$ always holds. If we include all the points in $X_{new} := \bigcup_{i=1,\dots,n} X_{new,i}$ in our set of test points $X_T(Q)$ (which is reasonable because in order to compute $\eta_{dir,i}$ we have to evaluate $\eta(x)$ for $x \in X_{new}$, anyway) then we can also ensure $\eta_{dir} \leq \tilde{\eta}_l$.

Now we refine the element only in those directions for which the corresponding test points yield large values, i.e., if the error estimate $\eta_{dir,1}$ is large we refine in x -direction and if the error estimate $\eta_{dir,2}$ is large we refine in y -directions (and, of course, we refine in both directions if all test points have large error estimates).

Anisotropic refinement can considerably increase the efficiency of the adaptive gridding strategy, in particular if the solution V has certain anisotropic properties, e.g., if V is linear or almost linear in one coordinate direction. Note that this is the case in our example. On the other hand, a very anisotropic grid Γ can cause degeneracy of the function V_Γ like, e.g., large Lipschitz constants or large (discrete) curvature even if V is regular, which might slow down the convergence. However, according to our numerical experience the positive effects of anisotropic grids are usually predominant.

References

- [1] Abel, A. (1990), Asset prices under habit formation and catching up with the Joneses, in: *American Economic Review* 40(2), 38-42.
- [2] Abel, A. (1999), Risk premia and term premia in general equilibrium, in: *Journal of Monetary Economics*, 43(1): 3-33.
- [3] Akdeniz, L. and W.D. Dechert (1997), Do CAPM results hold in a dynamic economy? *Journal of Economic Dynamics and Control* 21: 981-1003.
- [4] Alvarez-Cuadrado F., G. Monteiro and S.J. Turnovsky (2004), Habit formation, catching up with the Joneses, and economic growth, in: *Journal of Economic Growth*, 9, 47-80.
- [5] Bauer, F., L. Grüne and W. Semmler (2006), Adaptive spline interpolation for Hamilton-Jacobi-Bellman equations, *Applied Numerical Mathematics*, 56:1196-1210.
- [6] Becker, S., L. Grüne and W. Semmler (2006), Comparing second order approximation with dynamic programming, working paper, New School, New York, <http://www.newschool.edu/gf/cem>.
- [7] Breeden, D.T. (1979), An intertemporal asset pricing model with stochastic consumption and investment opportunities. *Journal of Financial Economics* 7: 231-262.
- [8] Boldrin, M., L.J. Christiano and J.D.M. Fisher (2001), Habit persistence, asset returns and the business cycle. *American Economic Review*, vol. 91, 1: 149-166.
- [9] Brock, W. (1979) An integration of stochastic growth theory and theory of finance, part I: the growth model, in: J. Green and J. Schenkman (eds.), New York, Academic Press: 165-190.
- [10] Brock, W. (1982) Asset pricing in a production economy, in: *The Economies of Information and Uncertainty*, ed. by J.J. McCall, Chicago, University of Chicago Press: 165-192.
- [11] Brock, W. and L. Mirman (1972), Optimal economic growth and uncertainty: the discounted case, *Journal of Economic Theory* 4: 479-513.
- [12] Burnside, C. (2001), Discrete state-space methods for the study of dynamic economies, in: Marimon, R. and A. Scott, eds., *Computational Methods for the Study of Dynamic Economies*, Oxford University Press, 95-113.
- [13] Candler, G.V. (2001), Finite-difference methods for continuous-time dynamic programming, in: Marimon, R. and A. Scott, eds., *Computational Methods for the Study of Dynamics Economies*, Oxford University Press, 172-194.
- [14] Campbell, J.Y. and J.H. Cochrane (1999), Explaining the poor performance of consumption-based asset pricing models, Working paper, Harvard University.
- [15] Chow, C.-S. and J.N. Tsitsiklis (1991), An optimal one-way multigrid algorithm for discrete-time stochastic control. *IEEE Trans. Autom. Control* 36:898-914.

- [16] Christiano, L. and J. Fisher (2000), Algorithms for solving dynamic models with occasionally binding constraints, *Journal of Economic Dynamics and Control* 24: 1179-1232.
- [17] Cochrane, J. (2001), *Asset pricing*, Princeton University Press, Princeton.
- [18] Collard, F. and M. Juillard (2001), Accuracy of stochastic perturbation methods: The case of asset pricing models; *Journal of Economic Dynamics and Control* 25: 979-999.
- [19] Constantinides, G.M. (1990), Habit formation: a resolution of the equity premium puzzle, in: *Journal of Political Economy* 98: 519-543.
- [20] Constantinides, G.M. and D. Duffie (1996), Asset pricing with heterogeneous consumers, in: *Journal of Political Economy* 104: 519-240.
- [21] Daniel, J.W. (1976), Splines and efficiency in dynamic programming, *J. Math. Anal. Appl.* 54:402–407.
- [22] Falcone, M. (1987), A numerical approach to the infinite horizon problem of deterministic control theory, *Appl. Math. Optim.*, 15: 1–13
- [23] Falcone, M. and R. Ferretti (1998), Convergence analysis for a class of high-order semi-Lagrangian advection schemes. *SIAM J. Numer. Anal.* 35:909–940.
- [24] Fuhrer, J.C. (2000), Habit formation in consumption and its implications for monetary-policy models, in: *American Economic Review* 90:367-390.
- [25] Grüne, L. (1997), An adaptive grid scheme for the discrete Hamilton-Jacobi-Bellman equation, *Numer. Math.*, 75: 1288-1314.
- [26] Grüne, L. (2003), Error estimation and adaptive discretization for the discrete stochastic Hamilton–Jacobi–Bellman equation. Preprint, University of Bayreuth. Submitted, <http://www.uni-bayreuth.de/departments/math/~lgruene/papers/>.
- [27] L. Grüne and W. Semmler (2004), Using dynamic programming with adaptive grid scheme for optimal control problems in economics. *Journal of Economic Dynamics and Control*, 28: 2427-2436.
- [28] Grüne, L. and W. Semmler (2004), Solving asset pricing models with stochastic dynamic programming, CEM working paper no. 54, Bielefeld University.
- [29] Grüne, L. and W. Semmler (2005), Asset pricing with loss aversion, working paper, New School, <http://www.newschool.edu/gf/cem>.
- [30] Hansen, L.P. and T. Sargent (2002) Robust control, book manuscript, Stanford University.
- [31] Jerman, U.J. (1998), Asset pricing in production economies, *Journal of Monetary Economics* 41: 257-275.

- [32] Johnson, S.A., J.R. Stedinger, C.A. Shoemaker, Y. Li, J.A. Tejada–Guibert (1993), Numerical solution of continuous–state dynamic programs using linear and spline interpolation, *Oper. Research* 41:484–500.
- [33] Judd, K.L. (1996), Approximation, perturbation, and projection methods in economic analysis, Chapter 12 in: Amman, H.M., D.A. Kendrick and J. Rust, eds., *Handbook of Computational Economics*, Elsevier, pp. 511–585.
- [34] Judd, K.L. and S.-M. Guu (1997), Asymptotic methods for aggregate growth models. *Journal of Economic Dynamics & Control* 21: 1025-1042.
- [35] Keane, M.P. and K.I. Wolpin (1994), The Solution and estimation of discrete choice dynamic programming models by simulation and interpolation: Monte Carlo evidence, *The Review of Economics & Statistics*, 76:648–672.
- [36] Kim, J. (2002), Indeterminacy and investment adjustment costs: an analytical result. Dept. of Economics, University of Virginia, mimeo.
- [37] Lettau, M. and H. Uhlig (1999), Volatility bounds and preferences: an analytical approach, revised from CEPR Discussion Paper No. 1678.
- [38] Lettau, M. and H. Uhlig (2000), Can habit formation be reconciled with business cycle facts? *Review of Economic Dynamics* 3: 79-99.
- [39] Lettau, M. G. Gong and W. Semmler (2001), Statistical estimation and moment evaluation of a stochastic growth model with asset market restrictions, *Journal of Economic Organization and Behavior*, Vol. 44: 85-103.
- [40] Lucas, R. Jr. (1978), Asset prices in an exchange economy. *Econometrica* 46: 1429-1446.
- [41] Marcet, A. (1994) Simulation analysis of stochastic dynamic models: applications to theory and estimation, in: C.A. Sims, ed., *Advances in Econometrics*, Sixth World Congress of the Econometric Society, Cambridge University Press, pp. 81–118.
- [42] McGrattan, E.R. and E.C. Prescott (2001) Taxes, regulation and asset prices. Working paper, Federal Reserve Bank of Minneapolis.
- [43] Munos, R. and A. Moore (2002), Variable resolution discretization in optimal control, *Machine Learning* 49:291–323.
- [44] Reiter, M. (1999), Solving higher–dimensional continuous–time stochastic control problems by value function regression, *J. Econ. Dyn. Control*, 23:1329–1353.
- [45] Rouwenhorst, K.G. (1995), Asset pricing implications of equilibrium business cycle models, in: T. Cooley: *Frontiers of Business Cycle Research*. Princeton, Princeton University Press: 295-330.
- [46] Rust, J. (1996), Numerical dynamic programming in economics, in: Amman, H.M., D.A. Kendrick and J. Rust, eds., *Handbook of Computational Economics*, Elsevier, pp. 620–729.

- [47] Rust, J. (1997), Using randomization to break the curse of dimensionality, *Econometrica* 65:478–516.
- [48] Ryder, Jr.H. and G. Heal (1973), Optimal growth with intertemporally dependent preferences, in: *Review of Economic Studies*, 40: 1-31.
- [49] Santos, M.S. and J. Vigo–Aguilar (1998), Analysis of a numerical dynamic programming algorithm applied to economic models, *Econometrica* 66: 409-426.
- [50] Schmidt-Grohe, S. and M. Uribe (2004), Solving dynamic general equilibrium models using a second-order approximation to the policy function, *Journal of Economic Dynamics and Control*, 28: 755-775.
- [51] Semmler, W. (2006), *Asset prices, booms and recessions - financial economics from a dynamic perspective*, Springer Publishing House, Heidelberg and New York, 2nd edition.
- [52] Tauchen, G. and R. Hussey (1991), Quadrature–based methods for obtaining approximate solutions to nonlinear asset–price models, *Econometrica*, 59:371–396.
- [53] Taylor J.B. and H. Uhlig (1990), Solving non-linear stochastic growth models: a comparison of alternative solution methods. *Journal of Business and Economic Studies*, 8: 1-18.
- [54] Trick, M.A. and S.E. Zin (1993), A linear programming approach to solving stochastic dynamic programs, Working Paper, Carnegie–Mellon University.
- [55] Trick, M.A. and S.E. Zin (1997), Spline approximations to value functions: a linear programming approach, *Macroeconomic Dynamics*, 1:255-277.