Computing equilibria in infinite-horizon finance economies: The case of one asset

Kenneth L. Judd[a,*], Felix Kubler[b], Karl Schmedders[c]

[a]*Hoover Institution, Stanford University, Stanford, CA 94305, USA*
[b]*Department of Economics, Stanford University, Stanford, CA 94305, USA*
[c]*Kellogg Graduate School of Management, Northwestern University, 2001 Sheridon Rd., Evanston, IL 60208, USA*

Abstract

We develop methods to compute equilibria in dynamic models with incomplete asset markets and heterogeneous agents. Using spline interpolation methods we approximate recursive trading policies of the agents and the equilibrium pricing functions. We explore various methods for determining the coefficients of these approximations, including time iteration methods and acceleration techniques. Exploring the optimization errors implied by the approximate equilibrium rules we examine the quality of our results. The results are very encouraging since we are able to compute approximate equilibria in a few minutes or less, attaining optimization errors on the order of one dollar per million dollars of wealth. © 2000 Elsevier Science B.V. All rights reserved.

E-mail address: judd@hoover.stanford.edu (K.L. Judd).

1. Introduction

The Lucas asset pricing model (Lucas, 1978) has had an important influence on financial economics and macroeconomics. However, the model assumes that there exists a representative agent. Consumer heterogeneity is an obvious fact of life and an appealing assumption from a theoretical point of view — with several agents the model constitutes a generalization of the Arrow-Debreu Model of general equilibrium under uncertainty (see Debreu, 1959) to a world with incomplete markets, an infinite time horizon, and smooth discounted utility functions. Moreover, it is well-known that the representative agent model fares poorly in explaining observed security prices (see for example Hansen and Singleton, 1982). The joint hypothesis of incomplete consumption insurance and consumer heterogeneity seems capable of enriching the pricing implications of the original model (see Constantinides and Duffie, 1996).

While there are now proofs of existence of an equilibrium in Lucas-type infinite-horizon exchange economies with heterogeneous agents and incomplete markets (see Magill and Quinzii (1996) or Hernandez and Santos (1996)), little is known about the quantitative nature of these equilibria. In general it is impossible to compute these equilibria because that would be equivalent to computing an infinite number of equilibrium prices. Duffie et al. (1994) show that if the exogenous shocks follow a finite-valued time-homogeneous Markov process then there exists an equilibrium in which the evolution of the endogenous variables follows a time-homogeneous Markov process having a time invariant transition with an ergodic measure.

When computing equilibria in infinite-horizon models, one wants to focus on dynamically simple, time-homogeneous Markov equilibria. The simplest such processes are recursive equilibria where the current endogenous variables are functions of the current state and where this state is of low dimension. Therefore, one wants to establish the existence of such equilibrium representations and approximate them by some finite parameterization.

There are several attempts in the literature to compute recursive equilibria in models with heterogeneous agents. Telmer (1993) considers a model with two agents and a single bond and assumes that exogenous shocks together with agents' current period bond holdings form a sufficient statistic for the future evolution of the economy. He discretizes the state space and uses a Gauss–Seidel strategy which searches for asset prices and an allocation of assets which comes as close as possible to clearing the market given the discretization of the state space. Because of the discrete state space this approach is likely to yield large approximation errors. Since in Telmer's model there is only a short-lived bond he avoids some of the technical problems. His errors are relatively small since for each exogenous shock he uses 1000 grid points to represent the endogenous state space. It seems however that his techniques cannot be extended to a more general model. (In his model he only considers three values for the exogenous

shock. One would imagine that running times rapidly increase with the number of shocks.) Heaton (1994), Heaton and Lucas (1996) and Zhang (1997) use similar techniques to compute equilibria. den Haan (1997) develops an algorithm to compute an equilibrium when there is a continuum of ex ante identical agents who differ in the realization of their ex post endowments. Also Krusell and Smith (1997) compute equilibria for models with a continuum of agents.

In this paper we offer an alternative approach for the case where there are two types of investors and one long-lived risky asset. We use spline collocation methods to solve for the approximations to the equilibrium price and asset investment policy functions. We use a time iteration method to determine the spline coefficients. In each iteration the algorithm takes as given the 'next' period's equilibrium functions and computes the 'current' period's equilibrium. To do this, we need only solve the agents' Euler equations together with the market clearing condition. Instead of searching for portfolio holdings over a grid space we use homotopy methods to solve the system of nonlinear equations under consideration. We then interpolate the resulting holdings and prices with cubic splines and use the new spline coefficients for the policy function in the next iteration. Coleman (1990) solves a stochastic growth model via policy function iteration using bilinear approximations of the functions. We use cubic splines, resulting in smoother approximations and smoother problems for our nonlinear equations solver.

Our solutions are approximations and depend only on the equilibrium behavior on the spline nodes we use. To determine the global quality of our solution, we examine the Euler equation errors at values of the state which are not used in our spline collocation procedure. We find that our maximum (relative) errors lie in the range of 10^{-6} and that by increasing the number of interpolation nodes we can reduce this error even further. This implies that our solution has agents making, in each period, optimization errors *less than one dollar per million dollars of wealth*. Our algorithm can handle general von Neumann–Morgenstern preferences, that is we allow agents to have different utility functions and different time discount factors. Using acceleration techniques our algorithm can compute equilibria for dynamic systems with 8 exogenous income states following a Markov process and calibrated to annual data in less than 20 s on a 233 MHz Pentium PC.

Previous work focussed on calibrations which corresponded to the period of time being one year. Since much of the interesting behavior in asset markets takes place at much higher frequency, it is important to be able to solve models corresponding to quarterly, monthly, and even weekly periods of time. Using spline collocation and acceleration techniques we can easily compute equilibria when the income process is calibrated to quarterly or monthly data and the agents' discount factors are assumed to be close to one. For a simple economy with i.i.d. income states we show how first and second moments of prices, returns

and trading volume vary significantly with different choices of the length of a time period.

The paper is organized as follows: In Section 2 we describe the economic model under consideration. Section 3 discusses conditions under which there exists a time-homogeneous equilibrium transition. Section 4 describes the algorithm and various technical issues. In Section 5 we illustrate the reliability of the algorithm and compute some examples to illustrate the algorithm. Section 6 reports running times for the algorithm and various accelerations methods.

2. The model

We consider a standard infinite-horizon pure exchange economy. Time is indexed by $t \in \mathbb{N}_0 \equiv \{0, 1, \dots\}$. A time-homogeneous Markov process of exogenous income states $(y_t)_{t \in \mathbb{N}_0}$ is valued in a discrete set $Y = \{1, 2, \dots, S\}$. The underlying probability space is denoted by $(\Omega, \mathcal{F}, Q)$ and the transition matrix by P. We assume that all elements of P are positive. Whenever we think of the income state in the next period as a random variable we denote it by $\tilde{y}$. A tribe $\mathcal{F}_t \subset \mathcal{F}$ generated by $\{y_0, \dots, y_t\}$ summarizes the information available at each time t. Finally, the filtration $F_t = \{\mathcal{F}_0, \dots, \mathcal{F}_t\}$ depicts how information is revealed through time t.

There are two types of infinitely lived agents indexed by $h = 1, 2$, and there is a single perishable consumption good in each state. Agent h's individual labor endowment in period t given income state $y \in Y$ is $e_t^h = e^h(y_t) \in \mathbb{R}_{++}$. Note that the function $e^i \colon Y \to \mathbb{R}_+$ depends on the exogenous income state alone. In order to transfer wealth across time and states agents trade in each period a long-lived asset paying a dividend $d \colon Y \to \mathbb{R}_+$ at an ex-dividend price q_t. The asset is in unit net supply. We denote agent h's security position in period t by $\theta_t^h \in \mathbb{R}$ and his initial endowment of the security by θ_{-1}. The aggregate endowment of the economy in period t is denoted by $e_t(y_t) = e_t^1(y_t) + e_t^2(y_t) + d_t(y_t)$.

Each agent h has von Neumann–Morgenstern preferences which are defined by a strictly monotone C^2, concave utility function $u_h \colon \mathbb{R}_{++} \to \mathbb{R}$ possessing the Inada property, that is, $\lim_{x \to 0} u_h'(x) = \infty$, and a discount factor $\beta_h \in (0, 1)$. For any $\mathcal{F}_t$-adapted consumption sequence $c = (c_0, c_1, c_2, \dots)$ the associated utility for agent h equals:

$$U_h(c) = \mathrm{E}\left\{\sum_{t=0}^{\infty} \beta_h^t u_h(c_t)\right\}.$$

Let

$$e = \begin{pmatrix} e^1(1) & e^2(1) \\ \vdots & \vdots \\ e^1(S) & e^2(S) \end{pmatrix}, \quad d = \begin{pmatrix} d(1) \\ \vdots \\ d(S) \end{pmatrix}$$

denote the matrix of possible individual endowments, $\boldsymbol{e}$, and the vector of possible dividends, d, respectively. Let $\mathscr{U}$ denote the set of all utility functions satisfying our assumptions above. The primitives of the economy can then be summarized as

$$\mathscr{E} = (\boldsymbol{e}, d, \boldsymbol{P}, \beta_1, \beta_2, u_1, u_2) \in \mathbb{R}_{++}^{2S+S+S\cdot S+2} \times \mathscr{U} \times \mathscr{U}.$$

2.1. Competitive equilibrium

An equilibrium is defined as a collection of $\mathscr{F}_t$-measurable portfolio holdings $\{(\theta_t^1, \theta_t^2)\}$ and asset prices $\{q_t\}$ such that security markets clear and all agents maximize their utility over their budget sets given the prices. This is a special case of the model considered in Magill and Quinzii (1996) or Hernandez and Santos (1996). In order to prove existence of an equilibrium one has to make an assumption on the agents' budget sets to rule out Ponzi Schemes (the indefinite postponement of debt). The conventional transversality condition from representative agents models cannot be used since with incomplete markets the expected present value of future wealth is not unambiguously defined. Magill and Quinzii (1996) define a transversality condition for incomplete markets and show that under weak assumptions this condition is equivalent to a debt constraint. To close the model we therefore impose an implicit debt constraint as an additional requirement of equilibrium: each agent's portfolio process $\{\theta^i\}$ is required to satisfy $\sup_t |\theta_t^i q_t|_\infty < \infty$, where for a random variable x, $|x|_\infty = \sup_{\omega \in \Omega} |x(\omega)|$. Note that this constraint does not constitute a market imperfection — it is just needed to ensure existence of a solution to the agents' optimization problem and in equilibrium it will never be binding. This point is emphasized in Magill and Quinzii (1996) and in Levine and Zame (1996).

Definition 1. A competitive equilibrium for an economy $\mathscr{E}$ is a collection of $\mathscr{F}_t$-measurable portfolio holdings $\{(\theta_t^1, \theta_t^2)\}$ and asset prices $\{q_t\}$ such that:
(1) $\theta^1 + \theta^2 = 1$, and
(2) for each agent i, $\theta^i \in \arg\max U_i(c)$ s.t. $c_t = e_t + \theta_{t-1}(q_t + d_t) - \theta_t q_t$ and $\sup_t |\theta_t^i q_t|_\infty < \infty$.

For a proof of existence see Magill and Quinzii (1996). Given our assumptions on the economies' fundamentals we can characterize equilibrium prices and portfolio holdings. While the implicit debt constraint is theoretically elegant it is difficult to impose it in a computational framework. Fortunately for the case of one infinitely lived asset there is an easy equivalent formulation of the constraint.

For each agent i define $\underline{\theta}^i = \max_{y \in Y} -(e^i(y)/d(y))$. Define an interval $I = (\underline{\theta}^1, 1 - \underline{\theta}^2)$.

Lemma 1. For each economy $\mathscr{E}$ agent 1's equilibrium portfolio holding has to lie in the bounded set I, that is in each equilibrium and for all t, $\theta_t^1 \in I$.

Proof. Assume that agent i's holding $\theta_t^i \leq \underline{\theta}^i$ at the end of period t. Let $y^* = \arg\max_{y \in Y} - (e^i(y)/d(y))$. Our assumption on $\boldsymbol{P}$ implies that for each time horizon T the probability of the next T income states being y^* is positive. Since agent i's consumption is restricted to be positive the agent has to decrease his asset holding every period in order to pay the dividends and to consume, that is $\theta_t^i > \theta_{t+1}^i$. This implies that with positive probability the agent's portfolio holding will violate any explicit debt constraint, and therefore also violates the implicit debt constraint. Because of market clearing agent 1's portfolio holding is also bounded from above. $\square$

Note that Lemma 1 implies that without loss of generality we can always impose a short-sale constraint. In equilibrium this constraint will never be binding.

For our discussion below it is important to emphasize that the agents' Euler equations together with market clearing fully characterize a competitive equilibrium.

Lemma 2. For an economy $\mathscr{E}$, $\mathscr{F}_t$-adapted processes $\{q_t\}$ and $\{(\theta_t^1, \theta_t^2)\}$ are an equilibrium if and only if for all t the following Euler equations and market clearing equation are satisfied with $c_t^i = e_t^i + \theta_{t-1}^i(q_t + d_t) - \theta_t^i q_t$:

$$q_t u_1'(c_t^1) - \beta_1 \mathrm{E}[(q_{t+1} + d_{t+1})u_1'(c_{t+1}^1)] = 0,$$

$$q_t u_2'(c_t^2) - \beta_2 \mathrm{E}[(q_{t+1} + d_{t+1})u_2'(c_{t+1}^2)] = 0,$$

$$\theta_t^1 + \theta_t^2 = 1.$$

Proof. See Constantinides and Duffie (1996, Appendix A). $\square$

3. Stationary Markov equilibria

In order to compute an equilibrium for an infinite horizon model it is necessary to focus on equilibria which are dynamically simple in the sense that one can choose a simple state space such that the current state is a sufficient statistic for the future evolution of the system and that this evolution can be approximated by a finite number of parameters. Because of agent heterogeneity the current exogenous state does not constitute such a sufficient statistic. The state space will also include endogenous variables, because some of them — such as for example the distribution of wealth or the agents' portfolio holdings — will clearly influence equilibrium prices. Duffie et al. (1994) examine

a model similar to ours and prove the existence of a time homogeneous Markov equilibrium (THME). They choose the endogenous state space to include current and last period portfolio holdings as well as current period prices and consumptions. Because in our model with two agents market clearing implies that agent 1's portfolio holding also describes agent 2's holding and because of Lemma 1 we can write the state space as $Y \times Z_{\text{THME}}$ with

$$Z_{\text{THME}} = \{(\theta_-, \theta, c, q) \in I \times I \times \mathbb{R}^2_{++} \times \mathbb{R}_{++}\}.$$

If for a set J, $\mathscr{P}(J)$ denotes the space of all measures on J, Duffie et al. (1994) show that there exists a measurable subset $J \subset Y \times Z_{\text{THME}}$ and a time invariant transition $\Pi: J \to \mathscr{P}(J)$, such that each time-homogeneous J-valued process with transition Π is an equilibrium for $\mathscr{E}$ and (J,Π) has an ergodic measure (Duffie et al., 1994, Theorem 3.1).

However, their approach of defining equilibrium as distributions is not very useful for computational purposes. While their approach has the advantage of being very general since the concept of an equilibrium defined as a distribution is very robust to taking limits (see Mas-Colell, 1992), it has the disadvantage that the existence theorem does not imply that Π satisfies any of the assumptions usually made in the applied literature such as, for example, the Feller property, i.e. that the transition function F is such that $\int_S h(s')F(s, ds')$ is continuous for a continuous function h (see for example Altug and Labadie (1994)). Without any notion of continuity of the transition Π it is not clear how to approximate it by a finite parameterization.

3.1. Recursive equilibria

The applied literature (see, e.g. Telmer (1993) or Heaton and Lucas (1996)) takes a somewhat different approach to this problem. The usual assumption made here is that the exogenous income state together with the individuals' portfolio holdings constitutes a sufficient minimal state space and that there exists a continuous policy function f as well as a function g mapping last period's portfolio holdings and current period's income state into current period prices. In our framework this would imply that the endogenous state space $Z = I$ and that the policy function $f: Y \times I \to I$ determines agent 1's optimal portfolio choice given portfolio holdings θ_- and the income state of the current period. Similarly, the price function $g: Y \times I \to \mathbb{R}_{++}$ maps agent 1's portfolio holdings θ_- and the income state of the current period into the asset price.

Let U be the set of all continuous functions $f: Y \times I \to I$ and let V be the set of all continuous $g: Y \times I \to \mathbb{R}_{++}$. Define

$$c^1(y, \theta_-) = e^1(y) + \theta_-(g(y, \theta_-) + d(y)) - f(y, \theta_-)g(y, \theta_-).$$

Let $W \subset U \times V$ be the set of all f and g such that $c^1(y, \theta_-) > 0$ and $e(y) - c^1(y, \theta_-) > 0$ for all $y \in Y$ and all $\theta_- \in I$. For all $y = 1, \ldots, S$ and $\theta_- \in I$ define a functional $F: W \rightarrow C \times C$ as follows:

$$F_1^y(f, g)(\theta_-) = g(y, \theta_-)u_1'(c) - \beta_1 E_y[(g_+ + d(\tilde{y}))u_1'(c_+)],$$

$$F_2^y(f, g)(\theta_-) = g(y, \theta_-)u_2'(e(y) - c) - \beta_2 E_y[(g_+ + d(\tilde{y}))u_2'(e(\tilde{y}) - c_+)],$$

where

$$c \equiv c^1(y, \theta_-), \quad c_+ \equiv c^1(\tilde{y}, f(y, \theta_-)), \quad g_+ \equiv g(\tilde{y}, f(y, \theta_-)).$$

We are now in the position to define the notion of a recursive equilibrium.

Definition 2. A recursive equilibrium for an economy $\mathscr{E}$ is a collection of functions $(f, g) \in W$ such that

$$F(f, g) = 0.$$

Lemma 2 immediately implies that every recursive equilibrium for $\mathscr{E}$ is in fact a competitive equilibrium as in Definition 1. Unfortunately, no general conditions on an economy's fundamentals are known which imply existence of a recursive equilibrium. Standard fixed point arguments cannot be applied to prove existence. In order to establish the existence of the functions f and g one would have to find conditions on the economy's fundamentals which make it possible to apply Schauder's theorem — a fixed point theorem for function spaces (see Zeidler (1986) for an overview of fixed-point theorems in infinite-dimensional spaces). This is beyond the scope of this paper.

The intuitive justification for assuming existence of a recursive equilibrium is that in each period with $y \in Y$ and initial portfolios $\theta_- \in I$ one can view the economy as a single-period economy with initial portfolios $\theta_- \in I$, income $y \in Y$, and specifications for the next period's equilibrium laws, f and g; let $\mathscr{E}(y, \theta_-; f, g)$ denote such an economy. From the results of Hernandez and Santos (1996, especially Corollary 3.7) we know that for all $y \in Y$ and all $\theta_- \in I$ there will be at least one competitive equilibrium for $\mathscr{E}(y, \theta_-; f, g)$. Define

$$(\bar{f}, \bar{g})(y, \theta_-; f, g) = \{(\theta_0, q_0): (\theta, q) \text{ is a competitive equilibrium}$$

$$\text{for } \mathscr{E}(y, \theta_-; f, g)\}.$$

Since the equilibrium prices and portfolio holdings will depend only on y and θ_- one might think that these two variables constitute a sufficient statistic for the evolution of the system.

However, this intuitive reasoning may fail and it seems plausible to argue that a recursive equilibrium does not always exist. Since generally the economy $\mathscr{E}$ will not have a unique equilibrium and since there are no known general conditions which guarantee uniqueness in this model, $(\bar{f}, \bar{g})$ will in general not be single

valued. This implies that it is not clear if there exists a single-valued selection (f, g) such that (f, g) constitutes a recursive equilibrium. First note that there is no reason why there should exist a continuous selection of the correspondence $(\bar{f}, \bar{g})$. One could imagine it to be similar to the equilibrium manifold for a simple Arrow–Debreu exchange economy. Such an equilibrium manifold is in general not convex-valued and there does not exist a continuous selection. Little is known about the equilibrium manifold for infinite horizon incomplete market models — if anything the correspondence $(\bar{f}, \bar{g})$ is likely to be more complex — for example it is not clear if equilibria are generically locally unique in this model (see Shannon, 1996).

Secondly, it is not clear whether there exists a selection at all (continuous or not) which satisfies $F(f, g) = 0$. Given an endogenous state $\theta_{t-1} \in I$ in period t there can be several prices and portfolio holdings satisfying the equilibrium conditions from t onwards while only one of them is compatible with the conditions up to period t. However at some period $s > t$ one might encounter the same endogenous state θ_{t-1} but different prices and portfolio holdings θ_s and q_s are compatible with the equilibrium conditions up to period s. It is clear that (y, θ_-) is not a sufficient statistic for the future evolution of the system.

3.2. Recursive equilibria and THME

The above discussion indicates that the larger state space Z_{THME} might be the right basis for defining a recursive equilibrium. In this case, given an endogenous state $z \in Z$ current period prices and portfolio holdings (q_t, θ_t) are compatible with all Euler equations if and only if they satisfy

1. $(q, \theta) \in (\bar{f}, \bar{g})(y_t, \theta_{t-1})$
2. For $c_t^i = e^i(y_t) + \theta_{t-1}^i(q_t + d_t) - \theta_t^i q_t,$

$$q_{t-1} u_1'(c_{t-1}^1) - \beta_1 \mathrm{E}[(q_t + d_t) u_1'(c_t^1)] = 0,$$

$$q_{t-1} u_2'(c_{t-1}^2) - \beta_2 \mathrm{E}[(q_t + d_t) u_2'(c_t^2)] = 0,$$

$$\theta_{t-1}^1 + \theta_{t-1}^2 = 1.$$

Applying Lemma 2 and with the above argument one obtains the following Conjecture.

Conjecture 1. *Given an economy $\mathscr{E}$ there exists a competitive equilibrium (q, θ) and a function* $f_{\text{THME}}: Y \times Y \times Z_{\text{THME}} \to Z_{\text{THME}}$, *such that for all* $t = 1, 2, \ldots,$

$$f_{\text{THME}}(y_{t+1}, y_t, \theta_{t-1}, \theta_t, c_t, q_t) = (\theta_t, \theta_{t+1}, c_{t+1}, q_{t+1}).$$

Since Duffie et al. (1994) are in general not able to rule out sunspots, they cannot prove the existence of a spotless Markov equilibrium. As they point out,

this implies that their result does not guarantee the existence of a policy function f_{THME} (see Duffie et al., 1994, Section 1.6).

Note that the conjecture — if it were true — does not imply that f_{THME} is continuous. Throughout the paper we will assume that f_{THME} is continuous. In order to prove continuity one has to apply Schauder's theorem, using a somewhat more complicated operator equation (in this case the Euler equations from two periods are used to describe equilibrium).

If one chooses the large state space Z_{THME} as the basis for a computational procedure one faces the problem that even with a single asset and two consumers one has five endogenous state variables. This causes various numerical problems. In order to avoid the problem of being faced with a high dimensional endogenous state space we make the following argument. Suppose a computational procedure finds an approximate recursive equilibrium, that is we find $(\hat{f}, \hat{g})$ such that for a small $\varepsilon > 0$ (say $\varepsilon = 10^{-8}$),

$$\sup_{\theta \in I} F(\hat{f}, \hat{g})(\theta) < \varepsilon.$$

Such as result trivially implies that we also computed an approximate f_{THME}, that is a policy function for a THME. A possible strategy for computing the policy function for a THME is therefore to assume first that there exists a recursive equilibrium (i.e. guess that f_{THME} does not depend on (c_t, q_t, θ_{t-1})) and to compute an approximate recursive equilibrium. If this procedure fails one has to use the larger state space. If it succeeds the initial guess was correct. In all the examples we considered the algorithm described below found an approximate recursive equilibrium, therefore we computed an approximate policy function for a THME. This argument leads to the following proposition.

Proposition 1. *There exists a recursive equilibrium for an economy $\mathscr{E}$ if there exists a competitive equilibrium with a continuous policy function f_{THME} for which*

$$\frac{\partial f(z)}{\partial \theta_-} = \frac{\partial f(z)}{\partial q} = \frac{\partial f(z)}{\partial c} = 0$$

for all $z \in Z$.

The fact that there existed an approximate recursive equilibrium for all the economies we considered indicates that the class of economies for which a recursive equilibrium exists is actually quite large. We want to argue that there are some nontrivial economies $\mathscr{E}$ for which a recursive equilibrium can be proven to exist.

It is well known that for economies where agents have identical CRRA utility functions and individual endowments are spanned by the securities' dividends there exists a Pareto-efficient equilibrium even when markets are incomplete. In

this equilibrium the exogenous income state alone constitutes a sufficient state (see Lucas, 1978). A recursive equilibrium as in Definition 2 exists.

Lemma 3. Let $\mathscr{E} = (\boldsymbol{e}, d, \boldsymbol{P}, \beta, \beta, u, u)$ be an economy with $u(c) = c^{1-\gamma}/1 - \gamma$ and with $e^1 = \xi_1 d$ and $e^2 = \xi_2 d$ for some $\xi_1, \xi_2 \in \mathbb{R}$. Then there exists a recursive equilibrium (f_0, g_0) for $\mathscr{E}$ with

$$f_0(y, \theta_-) = \theta_-$$

for all $\theta_- \in I$ and all $y \in Y$.

If endowments, dividends and preferences are 'close' to the representative agent case one would expect that a recursive equilibrium still exists. In order to verify this, one could possibly use the Implicit Function Theorem for Banach spaces (see, for example Zeidler, 1986, p. 151). The problem with this approach lies in the fact that (f_0, g_0) does not lie in the interior of W. It is subject to further research how to circumvent this problem and to generalize Lemma 3 to a much larger class of economies.

It is straightforward to extend our theoretical discussion of this section to the case of several infinitely lived assets and several agents. Since the emphasis of this paper lies on the computational aspects however, we chose not to introduce further assets or agents (see Judd et al., 1999, for the case of two assets and two agents).

4. The algorithm

The central theme of our algorithm is to approximate the policy functions f and g by cubic splines which we represent through B-splines and to compute the spline coefficients using collocation methods. We solve the collocation equations with an iterative approach.

4.1. Representing the equilibrium functions

It is necessary to globally approximate the functions f and g by finitely parameterized functions $\hat{f}, \hat{g}$ using relatively few parameters. Since monomials form a basis for the space of continuous functions a popular method of doing this is to represent the functions by a finite sum of orthogonal polynomials (see for example Judd (1992) who uses Chebychev polynomials and Judd (1998) for an overview over different approximation methods). However, functions exhibiting high curvature cannot be well approximated by orthogonal polynomials. Unfortunately, in our case the functions g frequently exhibit very high curvature near the boundaries of the interval I. It is for this reason that we chose to approximate the functions f and g by cubic splines. A cubic spline is a piecewise

polynomial function where the pieces are third-order polynomials and the function is twice continuously differentiable. It can be shown (see Judd (1998) for references) that cubic splines yield $O(n^{-4})$ convergence for $f \in C^4$. Since, unlike Chebychev polynomials, splines are good fits for functions with high curvature they perform much better than Chebychev polynomials in approximating the functions f and g. In the applications below the difference between the actually computed function and the values of the spline function were negligible.

For a general theory of splines and necessary programs for interpolation and computing B-splines see deBoor (1978). Here we only state the necessary information needed for our computations. B-splines of order 4 form a linearly independent basis for one-dimensional cubic splines. For our computations it is very helpful to represent the cubic splines by a sum of B-splines, since this representation reduces the number of free parameters by a factor of 4. Given a grid of knots (x_i) order k B-splines are recursively defined by

$$B_i^k(x) = \frac{x - x_i}{x_{i+k} - x_i} B_i^{k-1}(x) + \frac{x_{i+k+1} - x}{x_{i+k+1} - x_{i+1}} B_{i+1}^{k-1}(x)$$

with

$$B_i^0(x) = \begin{cases} 0, & x < x_i, \\ 1, & x_i \leq x \leq x_{i-1}, \\ 0, & x \geq x_i. \end{cases}$$

Given a function $f(x)$ and points $\{(x_1, f(x_1)), \ldots, (x_n, f(x_n))\}$ the knot sequence can be chosen such that there is a unique interpolating cubic spline $\hat{f}(x) = \sum_{i=1}^{n} \alpha_i B_i(x)$. The coefficients α_i can be obtained by solving a linear system of equations.

To compute the coefficients α we select a grid of as many mesh-points $M = \{\theta_j: j = 1, \ldots, n\}$ as we have unknown coefficients for each approximating function. For n mesh points we can write our approximating policy function $\hat{f}$ and price function $\hat{g}$ for an income state $y \in Y$ and a mesh-point $\theta_- \in M$ as follows:

$$\hat{f}(y, \theta_-) = \sum_{i=1}^{n} \alpha_{i,y}^f B_i(\theta_-)$$

and

$$\hat{g}(y, \theta_-) = \sum_{i=1}^{n} \alpha_{i,y}^g B_i(\theta_-).$$

To simplify our notation we write $\alpha \equiv (\alpha_{i,y}^f, \alpha_{i,y}^g)$ for the collection of all coefficients.

4.2. Collocation

Given an income state y and agent 1's last-period portfolio holding θ_-, let

$$\hat{c}^1 = e^1(y) + \theta_-(\hat{g}(y, \theta_-) + d(y)) - \hat{f}(y, \theta_-)\hat{g}(y, \theta_-)$$

denote our approximation for agent 1's current period consumption, let

$$\tilde{c}^1_+ = e^1(\tilde{y}) + \hat{f}(y, \theta_-)[\hat{g}(\tilde{y}, \hat{f}(y, \theta_-)) + d(\tilde{y})] - \hat{g}(\tilde{y}, \hat{f}(y, \theta_-))\hat{f}(\tilde{y}, \hat{f}(y, \theta_-))$$

denote the approximation for agent 1's random next period consumption. Let $\hat{c}^2 = e(y) - \hat{c}^1$ and let $\tilde{c}^2_+ = e(\tilde{y}) - \tilde{c}^1$ denote agent 2's consumptions. Substituting all approximations into the system of Euler equations we obtain a system of $2Sn$ equations, where for $y = 1, \ldots, S$ and $\theta_- \in M$ we have

$$\begin{aligned}
\hat{g}(y, \theta_-)u'_1(\hat{c}^1) &= \beta_1 E_y[(\hat{g}(\tilde{y}, \hat{f}(y, \theta_-)) + d(\tilde{y}))u'_1(\tilde{c}^1_+)], \\
\hat{g}(y, \theta_-)u'_2(\hat{c}^2) &= \beta_2 E_y[(\hat{g}(\tilde{y}, \hat{f}(y, \theta_-)) + d(\tilde{y}))u'_2(\tilde{c}^2_+)].
\end{aligned} \tag{1}$$

Notice that the problem has been transformed from finding functions f and g solving the Euler equations over the continuous state space to finding a zero of a large system of nonlinear equations that has the real coefficients α as unknowns. Note that for the system of equations to be well-defined it is necessary that $\hat{c}^i, \tilde{c}^i_+ > 0$ for $i = 1,2$. For a given economy $\mathscr{E}$, denote the set of coefficients α satisfying this inequality for all $\theta_- \in I$ by A.

Two theoretical questions of central importance are whether or under what conditions the system possesses a solution, and if the functions $\hat{f}$ and $\hat{g}$ converge to the true policy functions f and g as the number of mesh points tends to infinity. While it remains unclear if (1) always possesses a solution, we always found an approximate solution yielding a very small error in all the examples calculated below.

Our discussion of the existence problem in Section 3 above makes clear that for general economies one cannot prove convergence of the algorithm as the number of mesh-points tends to infinity. (For this one would have to use the larger state space Z_{THME} and find conditions which ensure continuity of f_{THME}.) However, for the case of Lemma 3 it is trivial to see that (1) has a solution for all finite n and that these functions are the true policy functions. Using the finite-dimensional version of Implicit Function Theorem one can show convergence of spline collocation methods at least for economies close to the representative agent case. For a more general global analysis the approach described in Petryshyn (1993) seems most promising. He develops the concept of A-proper operators and shows that if an operator is A-proper and every finite-dimensional approximation has a solution, the actual operator equation has a solution and the approximations converge to the true solution.

The most important practical question is how to solve the nonlinear system of equations (1). Since the system is very ill-conditioned (for the examples below the

condition numbers were all around 10^{10} even close to the true solution) conventional Newton methods cannot be used for reasonable starting points. Attempts to solve the system with homotopy methods seem quite promising and are subject of further research. However the system of equations tends to be very large. With 30 mesh-points and 12 states we have 720 equations. For the case with two assets and 8 states (see Judd et al., 1999) the number of equations lies around 7200, ruling out any algorithm which uses the Jacobian of the system.

We develop a time iteration algorithm which has the advantage of being robust to ill-conditioning and which can handle large systems. While it cannot be proven to be globally convergent it converged in all the examples calculated below.

4.3. A time-iteration algorithm

The basic intuition for our iterative approach is that at each iteration i, we take next period's policy functions as given and compute this period's portfolio holdings and prices which satisfy the Euler equation. Given functions $\hat{f}_i$ and $\hat{g}_i$ we obtain $\hat{f}_{i+1}$ and $\hat{g}_{i+1}$ by interpolating these portfolio holdings and prices.

More formally, given functions $\hat{f}, \hat{g}$ with coefficients α and the mesh $M = \{\theta_1, \ldots, \theta_n\}$ define for all $y = 1, \ldots, S$:

$$\Theta(y, \alpha) = \begin{pmatrix} \theta(y, \theta_1, \alpha) \\ \vdots \\ \theta(y, \theta_n, \alpha) \end{pmatrix}, \qquad Q(y, \alpha) = \begin{pmatrix} q(y, \theta_1, \alpha) \\ \vdots \\ q(y, \theta_n, \alpha) \end{pmatrix},$$

where $(\theta(y, \theta_-, \alpha), q(y, \theta_-, \alpha))$ solves the system of Euler equations for $\theta_- \in M$:

$$q u_1'(e_1(y) + \theta_-(q + d(y)) - \theta q)$$
$$= \beta_1 \mathrm{E}_y[(\hat{g}(\tilde{y}, \theta) + d(\tilde{y})) \cdot u_1'(e_1(\tilde{y}) + \theta(\hat{g}(\tilde{y}, \theta) + d(\tilde{y})) - \hat{g}(\tilde{y}, \theta)\hat{f}(\tilde{y}, \theta))],$$
$$q u_2'(e(y) - e_1(y) - \theta_-(q + d(y)) + \theta q)$$
$$= \beta_2 \mathrm{E}_y[(\hat{g}(\tilde{y}, \theta) + d(\tilde{y})) \cdot u_2'(e(\tilde{y}) - e_1(\tilde{y}) - \theta(\hat{g}(\tilde{y}, \theta) + d(\tilde{y}))$$
$$+ \hat{g}(\tilde{y}, \theta)\hat{f}(\tilde{y}, \theta))], \tag{2}$$

where $\hat{f}$ and $\hat{g}$ are our spline functions with coefficients α. We define a function $G: A \to A$ where $G(\alpha)$ is the set of spline coefficients which interpolates the equilibrium solutions $(\theta(y, \theta_-, \alpha), q(y, \theta_-, \alpha))$. We want to find an $\alpha \in A$ such that $G(\alpha) = \alpha$. Given a starting point α_0 our basic time iteration algorithm sets

$$\alpha_{i+1} = G(\alpha_i). \tag{3}$$

The algorithm terminates if

$$\max_{y \in Y, \, (\theta_j \in M)_{j=1, \ldots, n}} \{|\theta(y, \theta_j, \alpha_{i+1}) - \theta(y, \theta_j, \alpha_i)|, |q(y, \theta_j, \alpha_{i+1}) - q(y, \theta_j, \alpha_i)|\} < \varepsilon$$

for some small $\varepsilon > 0$.

This algorithm will converge if G is a contraction mapping. Locally we have a more concrete condition for convergence. Suppose that at a fixed point of G, α^*, G is Lipschitz and that the spectral radius of the Jacobian of G, $\rho(\partial_\alpha G(\alpha^*))$, is less than 1. Then for α_0 sufficiently close to α^* the iteration process (3) is convergent. While we are not able to prove that any of these two conditions hold for general economies $\mathscr{E}$ our algorithm also has a nice economic intuition. If we start with $\hat{f} = \hat{g} = 0$ we can interpret the ith iteration $\hat{f}_i(y, 0.5)$ and $\hat{g}_i(y, 0.5)$ as approximations for first period portfolio holdings and prices for a i-period finite horizon economy. Convergence of our iterative procedure then amounts to convergence of finite-horizon equilibria to a THME of the infinite-horizon model.

4.4. Solving the Euler equations

During any iteration for given functions $\hat{f}_i$ and $\hat{g}_i$ it is necessary to solve the nonlinear system of Euler equations (2) at all spline interpolation nodes $\theta_- \in M$ and each income state $y \in Y$. If we do not have a good starting point, Newton-method-based algorithms for solving (2) are not likely to perform well because they are not globally convergent and because the system of equations is not well-conditioned for many values of the endogenous variables. In this case we have to use homotopy methods to solve (2). The key insight for solving system (2) is that it is similar to the equilibrium conditions of the well-known General Equilibrium Model with Incomplete Markets (GEI Model). Therefore, in order to solve system (2) we can apply — with some modifications — the algorithm developed by Schmedders (1998,1999) for the GEI Model.

In order to apply Schmedders' algorithm we have to define a homotopy function. Denoting the homotopy parameter by λ our homotopy equations given functions $\hat{f}_i$ and $\hat{g}_i$ were as follows:

$$
\begin{aligned}
0 = {}& qu_1'(e_1(y) + \theta_-(q + d(y)) - \theta^1 q) \\
& - (\beta_1 \mathrm{E}_y[(\hat{g}_i(\tilde{y}, \theta^1) + d(\tilde{y}))u_1'(e_1(\tilde{y}) + \theta^1(\hat{g}_i[\tilde{y}, \theta^1] + d(\tilde{y})) \\
& - \hat{f}_i[\tilde{y}, \theta^1]\hat{g}_i[\tilde{y}, \theta^1])]), 0 = \lambda qu_2'(e(y) - e_1(y) - \theta_-(q + d(y)) + \theta^1 q) \\
& - \lambda(\beta_2 \mathrm{E}_y[(\hat{g}_i(\tilde{y}, \theta^1) + d(\tilde{y}))u_1'(e(\tilde{y}) - e_1(\tilde{y}) - \theta^1(\hat{g}_i[\tilde{y}, \theta^1] + d(\tilde{y})) \\
& + \hat{f}_i[\tilde{y}, \theta^1]\hat{g}_i[\tilde{y}, \theta^1])]) - (1 - \lambda)((\theta_- - \theta^1) \qquad (4)
\end{aligned}
$$

The term $-(1 - \lambda)(\theta^1 - \theta_-)$ is the derivative of the penalty function $-\frac{1}{2}(1 - \lambda)(\theta^1 - \theta_-)^2$. The penalty function is concave and therefore the Euler equations are still necessary and sufficient for optimality. For $\lambda = 1$ system (4) is equivalent to (2) and so a solution to (4) for $\lambda = 1$ yields the desired solution of the Euler equations. Since in our model there exists only one asset with nonnegative payoffs in all states we cannot have asset redundancy leading to discontinuous demand functions. Nevertheless the penalty approach is very

advantageous for us since it allows for a convenient choice of starting point when solving system (2). In comparison to Schmedders (1998,1999) we modified the penalty function so that for $t = 0$ the holding θ^1 at the unique starting point of the homotopy is $\theta^1 = \theta_-$. In order to determine the price at the starting point the first equation of system (4) needs to be solved for $\theta^1 = \theta_-$. Since the price q appears only in the first term and we use utility functions with easily invertible derivatives we can give a closed-form solution for q. Alternatively, we could solve this first equation numerically using a Newton method.

While the homotopy approach always finds a solution this method is time consuming. After about 3 to 4 time iterations, the difference between the old and new policy functions are rather small. If we take the values of the old policy functions $(\hat{f}(y,\theta_-), \hat{g}(y, \theta_-))$ as starting points for the Newton method a solution to (2) can often be found in a fraction of the time necessary for finding a solution with homotopy methods. This fact allows us to greatly speed up the algorithm by using a Newton algorithm for solving (2) after about 4 time iterations. If the Newton method fails to find a solution, the homotopy method is used.

The Inada-condition on agents' preferences which is needed in order to guarantee that the agents' first-order conditions are necessary and sufficient (negative consumption cannot be optimal) causes great numerical problems when consumption is close to zero. There is no guarantee that agents consumption is bounded away from zero — in fact in most cases it will not be.[1] Therefore we have to allow the interval I to be so large that for values near the boundary of I one of the agents is almost bankrupt and his consumption is close to zero. For numerical reasons it is then necessary to extend the utility function for negative consumption. We applied a quadratic extension of the utility functions for sufficiently small consumption values. Since with this quadratic extension there is no guarantee that consumption will be nonnegative we need a device to ensure that the optimal solution to the Euler equation never violates our debt constraint. We examined two alternatives. One way is to introduce a short-sale restriction on the asset. The homotopy algorithm is easily revised using the methods of Garcia and Zangwill (1981); see Judd (1998) for a description of this approach. While this method does yield nonnegative consumption it also slows down the algorithm since the number of equations we have to solve at each

---

[1] To see that individual agents' optimality does not bound consumption away from zero consider the following situation. There are two possible shocks in each period, each shock associated with a different return of the risky asset and there is an agent with constant relative risk aversion. The solution to his Euler equations will be such that his optimal consumption will increase by some fraction in the state where the return on the risky asset is high and decrease by some fraction in the other state. Assume that the same exogenous shock occurs T times and that for this shock the agent's consumption decreases. Clearly, since with positive probability there is no upper bound for T there is no positive lower bound for the agents' consumption. The probability that his consumption will reach zero is still zero however.

Table 1
Time iteration algorithm

Step 0:	Select an error tolerance ε for the stopping criterion, a mesh $M = \{\theta_1, \ldots, \theta_n\}$, and initial guesses $\hat{f}_0, \hat{g}_0$
Step 1:	Given functions $\hat{f}_i, \hat{g}_i$, $\forall \theta_- \in M$, $\forall y \in Y$ compute $\theta(y, \theta_-)$ and $q(y, \theta_-)$ either by solving the homotopy equations (4) for $\lambda = 1$ or (if a good starting point is known from previous iterations) by solving the system of equations (2) with Newton's method
Step 2:	Use the solutions from Step 1 for prices and portfolio trades at the points (y, θ_-) in $M \times Y$ to compute, via spline interpolation, the new approximations $\hat{f}_{i+1}, \hat{g}_{i+1}$
Step 3:	Check stopping criterion: If $\max_{y \in Y, \theta_- \in M}\{|\hat{f}_{i+1}(\theta_-) - \hat{f}_i(\theta_-)|, |\hat{g}_{i+1}(\theta_-) - \hat{g}_i(\theta_-)|\}$ $< \varepsilon$ then go to Step 4. Otherwise increase i by 1 and go to Step 1
Step 4:	The algorithm terminates. Set $\hat{f} = \hat{f}_{i+1}, \hat{g} = \hat{g}_{i+1}$

mesh point and for each state increases to five. Alternatively we discouraged agents to short the asset beyond a certain limit by introducing a very big penalty on asset short-holdings through a penalty function of the form $K \max\{(\underline{\theta} - \theta^i)^\kappa, 0\}$ for $\kappa \in \{2, 4\}$. This penalty can easily be incorporated into both the homotopy and the Newton algorithm and results essentially in the same solution as a short-sale constraint.

4.5. Summary of the algorithm

Our algorithm is summarized in Table 1.

4.5.1. Implementation in FORTRAN
We implemented our algorithm in FORTRAN on a Pentium 233 MHz PC. For the homotopy-path following in Step 1 of the algorithm we used the software package HOMPACK, a suite of FORTRAN 77 subroutines for solving systems of nonlinear equations using homotopy methods (Watson et al., 1987). From the three algorithms implemented in HOMPACK we selected the one that tracks the homotopy path by solving an ordinary differential equation (subroutine FIXPDF, Watson, 1979) since it is generally the most robust path-following method of HOMPACK. We provided the software with the starting point of the homotopy path and the Jacobian matrix of the homotopy. For the Newton-solver we used HYBRD, a suite of FORTRAN 77 subroutines for solving nonlinear equations using Powell's hybrid method (which is essentially a modification of Newton's method). For the interpolation and the evaluation of the B-splines subroutines from deBoor (1978) were used.

4.5.2. Interpolation on a bounded set
The interpolation of the policy functions $\hat{f}(y, \cdot): I \to I$ on the interval I could lead to problems with the simulation of the model. A typical interpolation

procedure does not take into account any bounds on the function being interpolated. For every grid point $\theta_- \in I$ the first-order conditions ensure that $\hat{f}(y, \theta_-) \in I$. However, the interpolation might result in the existence of points $\tilde{\theta} \in I$ such that $\hat{f}(y, \tilde{\theta}) \notin I$. Such a function would lead to an error in the simulation of the model, when the function $\hat{f}$ is effectively composed with itself many thousand times. Since the interpolation is only valid for the interval I, an evaluation of $\hat{f}$ at a point $\hat{f}(y, \tilde{\theta}) \notin I$ could generate huge simulation errors.

We have always checked our simulations for the occurrence of function values outside the set I. While such values are theoretically possible, we have so far never detected them in any of our examples. In particular in the models with utility penalties these penalties result in a very strong inward-pointing form of the policy functions once the penalty becomes sufficiently large.

4.6. Improving running times

For large discount factors β the price function $\hat{g}_i$ converges rather slowly. For $\beta = 0.996$, which seems a plausible discount factor for monthly data it often took over 3000 iterations to reach a stopping criterion of $\varepsilon = 10^{-6}$. Since it is desirable to compute solutions with such short time periods, we would like to find ways to accelerate convergence.

The standard way to improve running times is to alter the iterative scheme. We tried several acceleration techniques in order to improve running times. A general theory of acceleration techniques is only available for linear systems (see for example Axelson (1994) for an overview). For some special cases the results carry over to nonlinear systems when the unknown is 'close' to the true solution (see Ortega and Rheinboldt, 1970). The intuition is, that for α close to the solution α^* Taylor's theorem allows us to use only the Jacobian of G in approximating $G(\alpha)$. We discuss methods to accelerate linear schemes and examine if they can also be applied to our nonlinear problem. Here we present some acceleration methods used in the linear operator literature.

4.6.1. Stationary iterative methods

We first consider a first-order stationary method. A first-order stationary method uses only $G(\alpha_i)$ and α_i to determine α_{i+1}. Given an iterative scheme $\alpha_{i+1} = G(\alpha_i)$ one can increase its rate of convergence by using the first-order extrapolation method

$$\alpha_{i+1} = \omega G(\alpha_i) + (1 - \omega)\alpha_i$$

for an acceleration-factor $\omega \in (0, 2)$. For linear systems of the form $\alpha_{i+1} = J \cdot \alpha_i$ the optimal ω can be found as follows. Assume that all the eigenvalues of J are

real and let λ_1 denote the smallest eigenvalue, let λ_n denote the largest. The optimal ω is then

$$\omega^* = \frac{2}{2 - \lambda_1 - \lambda_n}.$$

It can be shown that an acceleration-factor $\omega \geq 2$ is never optimal (see Judd, 1998).

Determining the optimal acceleration factor ω for our nonlinear problem is difficult. First it is very costly to compute all eigenvalues of the Jacobian. For a simple example of only two states and 15 nodes this takes around 30 s. Since one has to compute the Jacobian using finite differences, computing the Jacobian for a system with n unknowns amounts to $n + 1$ evaluations of $G(\cdot)$. Secondly, for most of the examples the Jacobian of G has complex as well as real eigenvalues. In this case there is even for linear systems no general theory of how to determine the optimal ω.

We can do better than first-order methods. Intuitively, we suspect that the last iterate α_{i-1} might also contain useful information for determining α_{i+1}. Second-order methods try to use this information. A second-order iterative scheme has the form

$$\alpha_{i+1} = \omega G(\alpha_i) + (\tau - \omega)\alpha_i + (1 - \tau)\alpha_{i-1}.$$

For $\tau = 1$, a second-order scheme degenerates into a first-order method. For the linear case — assuming all eigenvalues are real — it is again well-known how to optimally determine the acceleration factors ω and τ (see Axelson, 1994) from the smallest and largest eigenvalues. Again, these results do not apply directly to for our nonlinear problem, but they do motivate us to try them here.

4.6.2. The Chebyshev iterative method

Stationary iterative methods are characterized by the fact that the acceleration factors remain constant over all iterations. The Chebyshev iterative method is a variable parameter version of the second-order stationary methods above. The Chebyshev second-order iterative scheme has the form

$$\alpha_{i+1} = \omega_i G(\alpha_i) + (\tau_i - \omega_i)\alpha_i + (1 - \tau_i)\alpha_{i-1}$$

where

$$\tau_i = \frac{a + b}{2}\omega_i, \quad \omega_i = \left(\frac{a + b}{2} - \left(\frac{b - a}{4}\right)^2 \omega_{i-1}\right)^{-1} \quad \text{and} \quad \omega_0 = 4/(a + b).$$

Let $\lambda_{\min}$ and $\lambda_{\max}$ be estimates of the smallest and largest eigenvalue of the Jacobian of G, respectively (for the more general case of complex eigenvalues see Axelson (1994)). Then we want to take $0 < a = 1 - \lambda_{\max}$ and $b = 1 - \lambda_{\min}$. For all the examples in the applied parts below we implemented the Chebyshev

method with $a = 0.05$ and $b = 1.4$. For the linear case it is well known (see Axelson, 1994) that the asymptotic rate of convergence for this method is the same as for the second-order stationary method with optimal parameters. However the Chebyshev method is much less sensitive to estimates of the extreme eigenvalues and also applies to the case of complex eigenvalues.

5. Examples

We now present the details of some examples.

5.1. Errors and their interpretation

When approximating policy and price functions we have to deal with computational errors. Stopping rules do not specify when the approximate equilibrium prices are close to the true equilibrium prices but instead when the difference between consecutive iterates is small. Therefore it remains unclear how close the computed prices and portfolios are to the actual equilibrium prices and portfolios. However, for a given approximation it is straightforward to compute the errors in the agents' first-order conditions. For all computations in this paper we compute the maximum relative errors in agents' Euler equations and decrease the stopping criterion if these errors turn out to be too high.

Unfortunately, low errors in agents' Euler equations do not give any indication of how close we are to an equilibrium either. This is a well known fact and there are various interpretations of this in the literature. Judd (1992) argues that it is not sensible to expect infinite precision from agents and that therefore the computed prices and allocations are likely to be a good description of the actual economic outcome. For this line of reasoning it is important to show how small the errors actually are. Without knowing the actual solution this is not unambiguously possible. Judd (1992) suggests to evaluate the Euler equations at the computed prices and allocation and compute the wealth equivalent of the Euler equation residual when projected in directions not used to compute the approximation. A small error here would be consistent with the interpretation of an approximate equilibrium in the sense that agents are close to rational.

Since for θ^1 close to $\overline{\theta^1}$ consumption for agent one is close to zero, our quadratic extension of the utility function becomes relevant (the same is true for agent two when $1 - \theta^1$ lies close to $\underline{\theta^2}$). It is clear that in this region the computed errors in the Euler equation are not reliable and generally much larger than outside this region. Since the probability that in equilibrium θ will ever reach this region is negligible, larger errors in the Euler equations here do not affect the reliability of the computed equilibrium functions. Note, however, that in situations with heterogeneous time discount factors the boundary of the interval might be important (see the example in Section 5.2). We performed

a second check on the method by computing equilibria for the case where we know the solution. For dividends, endowments and preferences satisfying the conditions in Lemma 3 we know $f(\theta)$ and $g(\theta)$ for all $\theta \in I$. When we computed the equilibrium functions for such an economy the maximum errors were on the order of 10^{-16}, that is negligible, even close to the boundary of the interval I. A lower theoretical bound on any error is clearly machine precision — around 10^{-16} on most computers.

Below we will always report the maximum error in wealth equivalents from the Euler equations, since we, of course, consider cases where we do not know the true equilibrium.

5.2. Introductory example

The purpose of the example in this section is to show the versatility of our algorithm and to obtain a first impression about what the approximated functions $\hat{f}$ and $\hat{g}$ can look like. Also we show how we use $\hat{f}$ and $\hat{g}$ to obtain interesting parameters describing the economy through simulations of the model.

We consider a model with $S = 2$ exogenous income states. The Markov transition probabilities are $p(1|1) = p(1|2) = p(2|2) = p(2|2) = 0.5$. The agents' functions and discount factors are as follows:

$$u_1(c) = 60c - 5c^2, \quad \beta_1 = 0.96,$$

$$u_2(c) = \log(c), \quad \beta_2 = 0.94.$$

The asset dividends and the agents' endowments and original portfolio holdings are as follows:

$$e^1 = \begin{pmatrix} 1.5 \\ 1.5 \end{pmatrix}, \quad e^2 = \begin{pmatrix} 2.0 \\ 1.0 \end{pmatrix}, \quad d = \begin{pmatrix} 1.0 \\ 1.0 \end{pmatrix}, \quad \theta^1_{-1} = \theta^2_{-1} = \frac{1}{2}.$$

Note that this economy does not satisfy our assumptions from Section 2.1. There is no guarantee that agent 1's debt constraint will never bind. However, as we point out below, the differences in time preferences lead to an equilibrium where agent 1 is always very long in the asset.

For the given data we have to determine $\hat{f}$ and $\hat{g}$ over the interval $I = (-1.5, 2)$. We chose 30 interpolation nodes $\theta_- \in M$ with 5 nodes very close to each the left and the right boundary of I (distributed on 10% of interval length at either side) and the remaining 20 nodes in uniform distances distributed over the remaining middle portion (80%) of I. For the termination criterion

$$\max_{\theta_- \in M, y \in Y} \{|\hat{f}_i(y, \theta_-) - \hat{f}_{i-1}(y, \theta_-)|, |\hat{g}_i(y, \theta_-) - \hat{g}_{i-1}(y, \theta_-)|\} < \varepsilon$$
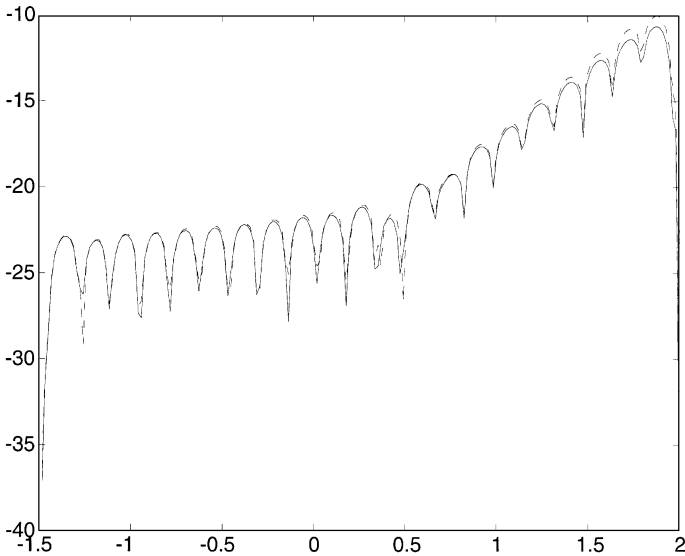
Fig. 1. Computational errors ($\log_{10}$) in Euler equations of Agent 1.

we choose $\varepsilon = 10^{-7}$. We should not accept $\hat{f}$ and $\hat{g}$ as acceptable approximations of the true functions f and g just because the termination criterion is satisfied. The algorithm which constructs $\hat{f}$ and $\hat{g}$ just checks the Euler equations at a small number of points. Before we accept $\hat{f}$ and $\hat{g}$, we should also check the quality of $\hat{f}$ and $\hat{g}$ at many other points in the state space. To check this, we plotted in Fig. 1 the base 10 logarithm of the absolute value of the relative errors in the Euler equation at various values of the state space which result from substituting $\hat{f}$ and $\hat{g}$ into the Euler equations. These errors represent the Euler equation error relative to the marginal utility of consumption, arriving at a relative wealth equivalent of the optimization errors of the agents. We plot the error function both income states, one represented by the broken line and the other by a solid line; the lines are barely distinguishable. The errors are consistently below 1×10^{-4} and get much lower as we move away from the right boundary, implying that the agents are making mistakes of less than one dollar per $10,000$ of consumption. The errors are relatively high for high θ_- because agent 2's Inada-condition causes numerical problems. Otherwise, the errors are quite small.

   Fig. 2 displays the portfolio trading policy computed by our algorithm, that is, the net investment functions $(\hat{f}(1,.) - \theta, \hat{f}(2,.) - \theta)$ of agent 1. The portfolio functions $\hat{f}$ reveal the striking differences between the two agents. Note that when agent 1 is very short in the asset, that is, when he is very poor, he still makes a large net investment into the asset. Obviously, he can only do this while allowing negative consumption. This is only possible because his quadratic utility function
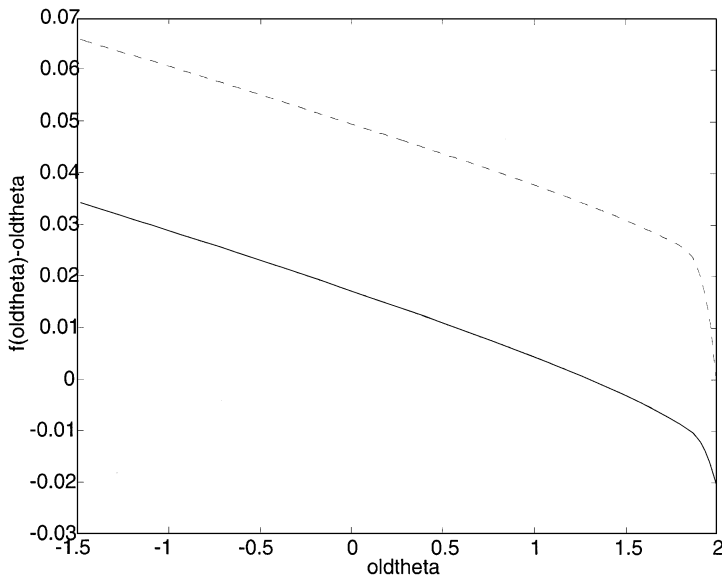
Fig. 2. Net investment functions of Agent 1 (oldtheta $\equiv \theta_-$).

violates the Inada condition. The behavior of agent 2 with his CRRA utility function satisfying the Inada condition is very different. When he is poor he barely changes his portfolio holdings. Note that as θ_- tends to $1 - \underline{\theta^2}$ agent 2's net investment is nonpositive in both states. This is guaranteed by the Inada condition.

We also find that the price of the asset depends crucially on the distribution of wealth in this example. When the agent with the high β owns most of the asset its price will be much (up to 50%) higher than when agent 2 owns it. A large variation in the distribution of portfolio holdings will lead to a high price volatility. The question is, however, how much variation of the distribution of portfolio holdings one can expect in equilibrium.

The next step of our analysis was to simulate the economy over 1500 time periods using the functions $\hat{f}$ and $\hat{g}$ and to determine the values of some interesting parameters. We repeated this simulation process 200 times. For each simulation of the economy we computed the average price and price variance as well as the average volume and the variance of the volume. The average of these values over the 200 runs were:

average price = 24.01,

price variance = 13.43,

average volume = 1.507×10^{-2},

volume variance = 2.607×10^{-2}.

5.3. Security prices in a stylized economy

To demonstrate how market incompleteness affects the behavior of aggregate time series we first examine a simple example. There are only two exogenous income states, there is no aggregate risk and a console in zero net supply paying one unit in each state and each period. The two agents have identical constant relative risk aversion and identical discount factors of $\beta = 0.99$. We considered two values for the coefficient of relative risk aversion, namely 1 and 4. The agents' endowments are

$$e^1 = \begin{pmatrix} 2.0 \\ 1.0 \end{pmatrix}, \quad e^2 = \begin{pmatrix} 1.0 \\ 2.0 \end{pmatrix}, \quad \theta^1_{-1} = \theta^2_{-1} = 0.$$

with the Markov transition probabilities $p(1|1) = p(2|2) = 0.9$ and $p(1|2) = p(2|1) = 0.1$.

It is clear that in the complete market solution c^1 and c^2 are constant over time and that the price of the console is $q_t = \beta/(1 - \beta) = 99$ for all t. However, when there is only one asset markets are incomplete and both the price and individual consumption vary over states and times. In Fig. 3 we display a
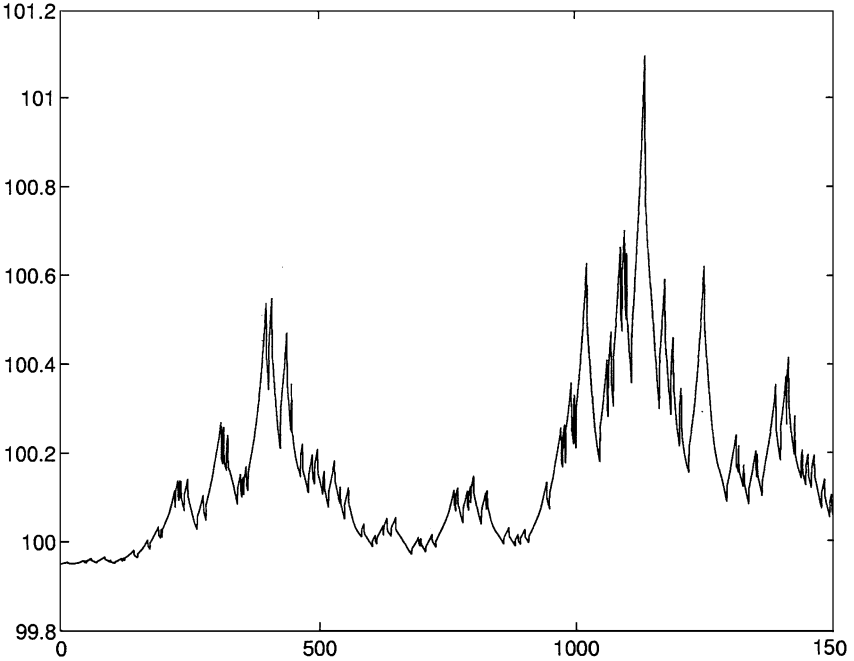

Fig. 3. Asset prices for $\gamma = 1$.

Table 2
Asset prices

		Average price	Volatility
Complete markets		99.0	0.0
Incomplete markets	$\gamma = 1$	100.40	0.32
	$\gamma = 4$	109.75	2.71

simulation of the price of the console over 1500 time periods for $\gamma = 1$. It is clear that the average price of the console is above 99 and that it varies significantly over time. The picture for $\gamma = 4$ is similar except that the price range is 108 to 118 instead of 100 to 102. This shows that the effects of incomplete markets are greatly magnified by moderately larger coefficients of relative risk aversion. In Table 2 we show the first and second moments of the price for the cases $\gamma = 1,4$. When agents have higher risk aversions the price of the asset is significantly larger and also more volatile.

5.4. Different time horizons

Earlier papers that computed dynamic stochastic equilibria with heterogeneous agents and incomplete markets calibrated their models to yearly data. The main reason for this is that when the time discount factor β lies close to 1 the system becomes quite ill-conditioned and one faces considerable numerical problems. Although even for the case of one asset these problems are severe, our methods allow us to compute equilibria for quarterly data ($\beta = 0.988$) and even monthly data ($\beta = 0.996$). We assume that each period there are shocks to each agents' labor income as well as to dividends and that all three shocks are independent and i.i.d. over time. The first question that arises is how to calibrate the shocks to the US economy for different time periods. A natural approach would be to choose the shocks such that first and second moments of total yearly dividends and labor incomes remain the same regardless of the length of the trading period. If the monthly dividend is $d \pm \varepsilon$ the dividend for n months then must be $nd \pm \varepsilon\sqrt{n}$. The formula for labor income is analogous. We considered three cases corresponding to annual, quarterly, and monthly calibrations.

Yearly calibration:

$$e^1 = \begin{pmatrix} 12 - \sqrt{0.12} \\ 12 + \sqrt{0.12} \end{pmatrix}, \quad e^2 = \begin{pmatrix} 18 - \sqrt{0.48} \\ 18 + \sqrt{0.48} \end{pmatrix}, \quad d = \begin{pmatrix} 7.8 - \sqrt{0.27} \\ 7.8 + \sqrt{0.27} \end{pmatrix},$$

$$\beta = 0.953.$$

Table 3
Return and trading volume for different time horizons

| β | 0.996 | 0.988 | 0.953 |
Time unit	month	quarter	yr
Average return per period	1.0058	1.0177	1.0508
Std. Dev. return per period	0.1875	0.1079	0.05525
Average trading volume per period	5.213×10^{-4}	8.989×10^{-4}	17.569×10^{-4}
Average annual return	1.0725	1.0610	1.0508
Average annual trading volume	6.421×10^{-3}	3.307×10^{-3}	1.757×10^{-3}

Quarterly calibration:

$$e^1 = \begin{pmatrix} 3 - \sqrt{0.03} \\ 3 + \sqrt{0.03} \end{pmatrix}, \quad e^2 = \begin{pmatrix} 4.5 - \sqrt{0.12} \\ 4.5 + \sqrt{0.12} \end{pmatrix}, \quad d = \begin{pmatrix} 1.95 - \sqrt{0.0675} \\ 1.95 + \sqrt{0.0675} \end{pmatrix},$$

$$\beta = 0.988.$$

Monthly calibration:

$$e^1 = \begin{pmatrix} 1 - 0.1 \\ 1 + 0.1 \end{pmatrix}, \quad e^2 = \begin{pmatrix} 1.5 - 0.2 \\ 1.5 + 0.2 \end{pmatrix}, \quad d = \begin{pmatrix} 0.65 - 0.15 \\ 0.65 + 0.15 \end{pmatrix},$$

$$\beta = 0.996.$$

Agents have identical constant relative risk aversion preferences with a coefficient of relative risk aversion of 1.5. Table 3 shows first and second moments of returns as well as trading volume for the three cases. We can make the following observations. Average returns are higher when we calibrate the economy to shorter time periods. Note that there is a negative serial correlation in returns — total yearly returns of the risky asset are not equal to the product of twelve monthly or four quarterly rates of return, respectively. These results contradict the intuition that the equilibrium yearly return of the risky security should only depend on the mean and the variance of dividends and income because the resolution of uncertainty is irrelevant. The yearly trading volume is larger when the trading periods are shorter. This result confirms the intuition that risk-sharing opportunities increase as the length of a trading period decreases.

6. Performance of the algorithm

In this section we discuss the performance of our algorithm and the various acceleration strategies. There are several factors which influence running times.

It is obvious from our iterative algorithm that the number of spline colloca-
tion points is almost proportionally related to running times. In most of
the applications below we used 15 collocation points. We placed 5 points
close to each boundary, where the functions f and g tend to exhibit high
curvature and 5 points in the middle. For less than 15 collocation points there
is a significant trade-off between running times and maximum errors in the
Euler equation. While the maximum errors can be reduced by using more than
15 points these gains seemed rather trivial.

The choice of the starting point, i.e. $\hat{f}_0$, $\hat{g}_0$ is another important determi-
nant of running times. When one wants to compute several similar examples
it is useful to use the computed policy functions of the previous example as
a new starting point. However, in order to give an objective evaluation of
the performance of the algorithm it seems reasonable to choose starting
points which are independent of the actual economy under considera-
tion — the reported running times below all refer to the case where $\hat{f}_0 \equiv 0$,
$\hat{g}_0 \equiv 0$.

6.1. Parameter specifications

We consider an economy with random dividends and idiosyncratic labor
incomes for each agent type. We assume that agents have identical utility
functions equal to $\log(c)$ and an identical discount factor β.

We consider the three cases of endowment and dividend processes. For Case
1 we assume that these three random variables are independent and that
they take two possible values with identical probabilities. We therefore have
8 income states,

$$e^1 = (1.9, 1.9, 1.9, 1.9, 2.1, 2.1, 2.1, 2.1),$$

$$e^2 = (1.8, 1.8, 2.2, 2.2, 1.8, 1.8, 2.2, 2.2),$$

$$d = (0.8, 1.2, 0.8, 1.2, 0.8, 1.2, 0.8, 1.2),$$

with each state s having probability $p(s) = 0.125$ independently of the previous
period's state. We also assume $\theta^1_{-1} = \theta^2_{-1} = \frac{1}{2}$.

For Case 2 there is in addition to the shocks in Case 1 a ninth exogenous
shock with very low probability in which agent 1 experiences a large
decrease in income. We choose $e^1_9 = 0$, $e^2_9 = 2.2$ and $d_9 = 0.4$. The transition
probabilities are $p(i|j) = 0.12375$, $p(i|9) = 0.05$, $p(9|i) = 0.01$ for $i, j < 9$
and $p(9|9) = 0.6$. In this case there is a precautionary motive for savings since,
in order to avoid starvation in state nine, agent one has to hold the risky
asset.

Table 4
Cases 1 and 2 running times with acceleration

(ω,τ)	Case 1			Case 2	
	$\beta = 0.95$	$\beta = 0.99$	$\beta = 0.996$	$\beta = 0.95$	$\beta = 0.99$
(1,1)	0:49	3:31	8:45	0:59	8:01
(1.25,1)	0:43	3:13	7:23	0:41	7:19
(1.5,1)	0:39	2:51	6:54	0:33	6:37
(1.75,1)	0:30	2:07	5:04	NC	NC
(1.75,1.6)	0:16	0:59	2:32	NC	NC

Table 5
Case 3 Running times with acceleration

$\omega\backslash\tau$	1.0	1.1	1.2	1.4	1.5	1.6
1.0	4:09	–	–	–	–	–
1.25	3:26	3:11	–	–	–	–
1.5	2:54	2:39	2:24	–	–	–
1.75	2:31	2:19	2:05	1:38	1:19	1:12
1.8	NC	2:15	2:02	–	–	–
1.9	NC	2:09	1:57	–	–	–
2.0	NC	NC	1:51	1:26	1:06	1:07
2.1	NC	NC	1:46	1:22	1:01	1:05
2.15	NC	NC	NC	1:20	1:00	1:02
2.25	–	–	–	1:16	0:57	NC
2.55	–	–	–	1:02	0:56	NC
2.6	–	–	–	NC	NC	NC

Case 3 is a variation of Case 2 with $\beta = 0.99$. There are nine states but the ninth state is not a bad state as in Case 2 but an average state.

6.2. Stationary iterative methods

We first consider a first-order stationary method with $\alpha_{i+1} = \omega G(\alpha_i) + (1 - \omega)\alpha_i$. Table 4 shows how different acceleration factors influence running times. The term NC refers to the case where the algorithm did not converge. Our stopping criterion was $\varepsilon = 10^{-6}$. The corresponding maximum errors in the Euler equations lie in the range of 5×10^{-7}. Table 4 indicates that for the nonlinear case there is not one 'optimal' acceleration factor as there is in the linear case. For the examples from Case 1, picking an acceleration factor between 1.5 and 1.75 seems like a good general strategy to improve convergence.

Table 6
Sequence of weights

Iter. #	τ_i	ω_i
1	1.7650	2.4344
2	1.6194	2.2336
3	1.5407	2.1250
4	1.5012	2.0706
5	1.4822	2.0444
10	1.4656	2.0216
15	1.4653	2.0211
20	1.4653	2.0211

However, as Case 2 shows this can be a dangerous strategy. Table 4 also shows that running times can be substantially improved by using second-order methods as opposed to first-order methods. Tables 5 displays the impact of second-order iterative accelerations applied to Case 3. Here, blank boxes indicate (ω, τ) which were not run.

In general it seems that first-order methods (that is, $\tau = 1$) do not improve running times dramatically. While second-order methods often yield much larger improvements in running times they are sensitive to the choice of the acceleration parameters. Second-order methods are recommended if one wants to compute several examples with similar endowments. In such cases it is worth the effort to try out several different (ω, τ) choices to find good combinations which can then be used across many different examples. Second-order methods are also recommended when discount factors are close to one since standard acceleration methods are quite slow.

6.3. The Chebyshev iterative method

We next consider the second-order Chebyshev iterative scheme with variable parameter. Table 6 shows how the values of τ_i and ω_i change over the iterations for estimates $a = 1 - \lambda_{\max} = 0.05$ and $b = 1 - \lambda_{\min} = 1.4$. Note that the sequences converge very fast and after 15 to 20 iterations the values barely change. The values are comparable to good values for the second-order stationary method.

As we pointed out above it does not make sense to determine the extreme eigenvalues for our complicated nonlinear problem. Table 7 shows how different estimates for a and b influence running times.

The table shows that for conservative estimates of the extreme eigenvalues the Chebyshev iterative method is very robust and yields considerable

Table 7
Running times depending on (a, b)

(a, b)		(0.05, 1.4)	(0.05, 1.6)	(0.02, 1.4)
Case 1a	$\beta = 0.99$	1:00	1:22	0:37
Case 1b	$\beta = 0.996$	2:31	3:04	1:41
Case 2	$\beta = 0.99$	2:07	2:39	NC
Case 3	$\beta = 0.99$	1:08	1:31	0:49

improvements of running times. The main advantage of the Chebyshev iterative method is that it is not very sensitive to the quality of the estimates of the extreme eigenvalues. This method should only be applied after a number of iterations of the algorithm and some significant progress towards the solution has been made.

7. Conclusion

This paper develops an algorithm to compute stationary equilibria in infinite horizon incomplete market models with heterogeneous agents. For the special case of two agents and only one asset we show that there exist stationary equilibria which can be described by a time-homogeneous policy function. Assuming that this function only depends on the agents' portfolio holdings and the exogenous shock, we develop a spline collocation algorithm to approximate the policy function. The algorithm is faster than previous methods, allowing us to compute solutions for models with short time periods and to do so with small error. In particular, in many models calibrated to quarterly data we can compute equilibria in a minute with optimization errors on the order of 10^{-6}. Errors of this magnitude make it seem likely that, at least for the cases we considered, we are justified to assume that the agents' portfolio holdings form a sufficient endogenous state space. We then show how these results can be used to derive results concerning asset pricing, trading volume, and investor welfare.

The algorithm can also be used to compute equilibria for models with several assets. Judd et al. (1999) consider a model where there is a short-term bond and an infinitely lived stock. While in this model it is not clear if there exists a recursive equilibrium, their results are also very encouraging. The spline collocation method combined with a time iteration process is a stable, fast, and accurate approach to dynamic recursive models.

References

Axelson, O., 1994. Iterative Solution Methods. Cambridge University Press, Cambridge.

Altug, S., Labadie, P., 1994. Dynamic Choice and Asset Markets. Academic Press, New York.

deBoor, C., 1978. A Practical Guide to Splines. Springer, Berlin.

Coleman, W.J., 1990. Solving the stochastic growth model by policy function iteration. Journal of Business and Economic Statistics 8, 27–29.

Constantinides, G.M., Duffie, D., 1996. Asset pricing with heterogeneous consumers. Journal of Political Economy 104, 219–240.

den Haan, W., 1997. Solving dynamic models with aggregate shocks and heterogeneous agents. Macroeconomic Dynamics 1, 355–386.

Debreu, G., 1959. A Theory of Value. Cowles Foundation Monograph, Yale University Press.

Duffie, D., Geanakoplos, J., Mas-Colell, A., McLennan, A., 1994. Stationary Markov equilibria. Econometrica 62, 745–781.

Garcia, C., Zangwill, W., 1981. Pathways to Solutions, Fixed Points, and Equilibria. Prentice-Hall, Englewood Cliffs, NJ.

Hansen, L.P., Singleton, K.J., 1982. Generalized instrumental variables estimation of nonlinear rational expectations models. Econometrica 50, 1269–1286.

Heaton, J., Lucas, D.J., 1996. Evaluating the effects of incomplete markets on risk sharing and asset pricing. Journal of Political Economy 104, 443–487.

Hernandez, A., Santos, M.S., 1996. Competitive equilibria for infinite-horizon economies with incomplete markets. Journal of Economic Theory 71, 102–130.

Judd, K.L., 1992. Projection methods for solving aggregate growth models. Journal of Economic Theory 58, 410–452.

Judd, K.L., 1998. Numerical Methods in Economics. MIT Press, Cambridge, MA.

Judd, K.L., Kubler, F., Schmedders, K., 1999. Incomplete asset markets with heterogeneous tastes and idiosynchratic income. Hoover Institution working paper.

Krusell, P., Smith, A.A., 1997. Income and wealth heterogeneity, portfolio choice, and equilibrium asset returns. Macroeconomic Dynamics 1, 387–422.

Levine, D., Zame, W., 1996. Debt constraint and equilibrium in infinite horizon economies with incomplete markets. Journal of Mathematical Economics 26, 103–131.

Lucas, R.E., 1978. Asset prices in an exchange economy. Econometrica 46, 1429–1445.

Magill, M., Quinzii, M., 1996. Incomplete markets over an infinite horizon: long-lived securities and speculative bubbles. Journal of Mathematical Economics 26, 133–170.

Mas-Colell, A., 1992. Infinite dimensional equilibrium theory: discussion of Jones. In: Laffont, J.J. (Ed.), Advances in Economic Theory, Sixth World Congress. Cambridge University Press, Cambridge.

Ortega, J., Rheinboldt, W., 1970. Iterative Solution of Nonlinear Equations in Several Variables. Academic Press, New York.

Petryshyn, W.V., 1993. Approximate-Solvability of Non-Linear Functional and Differential Equations. Marcel Dekker, New York.

Schmedders, K., 1998. Computing equilibria in the general equilibrium model with incomplete asset markets. Journal of Economic Dynamics and Control 22, 1375–1401.

Schmedders, K., 1999. A homotopy algorithm and an index theorem for the general equilibrium model with incomplete asset markets. Journal of Mathematical Economics, forthcoming.

Shannon, C., 1996. Determinacy of competitive equilibrium in economies with many commodities. UC Berkeley working paper.

Telmer, C.I., 1993. Asset-pricing puzzles in incomplete markets. Journal of Finance 48, 1803–1832.

Watson, L.T., 1979. A globally convergent algorithm for computing fixed points of C2 maps. Applied Mathematics and Computation 5, 297–311.

Watson, L.T., Billups, S.C., Morgan, A.P., 1987. HOMPACK: A suite of codes for globally convergent homotopy algorithm. ACM Transactions on Mathematical Software 13, 281–310.

Zeidler, E., 1986. Nonlinear Functional Analysis and Its Applications I. Springer, New York.

Zhang, H.H., 1997. Endogenous short-sale constraint, stock prices and output cycles. Macro-economic Dynamics 1, 228–254.