Journal of Economic Dynamics & Control 34 (2010) 28–35

Contents lists available at ScienceDirect

Journal of Economic Dynamics & Control

journal homepage: www.elsevier.com/locate/jedc

Solving the incomplete markets model with aggregate uncertainty by backward induction

Michael Reiter *

Institute for Advanced Studies, Stumpergasse 56, A-1060 Vienna, Austria

ARTICLE INFO

Article history:
Received 17 September 2008
Accepted 26 November 2008
Available online 8 August 2009

JEL classification:
C63
C68
E21

Keywords:
Heterogeneous agents
Backward induction

ABSTRACT

This paper describes a method to solve models with a continuum of agents, incomplete markets and aggregate uncertainty. I use backward induction on a finite grid of points in the aggregate state space. The aggregate state includes a small number of statistics (moments) of the cross-sectional distribution of capital. For any given set of moments, agents use a specific cross-sectional distribution, called "proxy distribution", to compute the equilibrium. Information from the steady state distribution as well as from simulations can be used to chose a suitable proxy distribution.

1. Introduction

This paper describes a method to solve the model with incomplete markets and aggregate uncertainty as specified in Den Haan (2009). The method is applicable to a large class of models with a continuum of agents and incomplete markets.

Early solution methods (Den Haan, 1996; Krusell and Smith, 1998) iterate between solving the household problem, conditional on an assumed aggregate law of motion, and updating the aggregate law of motion, based on the policy function of the household. Iteration continues until consistency is achieved between the solution of the individual household problem and the aggregate law of motion. The idea of the method I present here is to solve the model in one run of backward iterations, on a discrete grid of points in the aggregate state space. Consistency between individual and aggregate solution is enforced in each step of the backward iteration. This can be done separately point for point on the grid of aggregate states.

Following Den Haan (1996) and Krusell and Smith (1998), I work with the approximation that household decisions are based not on the complete state of the model, which is infinite-dimensional because it contains the whole cross-sectional distribution, but rather on an n_m-vector of statistics (or "moments") m, which characterize this distribution. This gives rise to a major complication: in order to check whether at the point in the state space characterized by m, individual behavior is consistent with the aggregate law of motion, it is not sufficient to know the statistics m, but we need to know the whole cross-sectional distribution. We therefore need a selection mechanism that tells us which distribution to use when

* Tel.: +43 1 59991 154.
E-mail address: michael.reiter@ihs.ac.at

0165-1889/$ - see front matter
doi:10.1016/j.jedc.2008.11.009

statistics are equal to m. I call this a "distribution selection function" (DSF), a mapping from aggregate states to cross-sectional distributions:

$$DSF: \quad (m,a) \to F(k,\varepsilon;m,a) \tag{1}$$

Here, a is the exogenous state variable (aggregate productivity). The DSF provides a family of distributions, parameterized by the statistics m. I call the distribution $F(k,\varepsilon;m,a)$ selected by (1) the "proxy distribution"; it stands in for the true cross-sectional distribution, which the household does not know or is not supposed to utilize.

The advantages of this method are

- It is fast, because it does not need the repeated simulation and solution of the model. This becomes more important if more than one or two statistics are used to describe the distribution.
- It depends very little (in simple cases not at all) on simulations of the model, thereby avoiding problems of sampling error that come with simulations.
- It does not rely on a parameterization of the aggregate transition law. At any point m in the state space we solve for next period's statistics m' (conditional on the aggregate shock).

The disadvantage is that the method relies on a DSF. In Krusell and Smith (1998), the distributions are generated through simulation, which forms part of the solution procedure.

2. The method

The outline of the algorithm is the following:

(1) Choose a discrete representation of the economy. This includes
  - a finite parameterization of the cross-sectional distribution function (not to be confounded with the vector of statistics m);
  - a finite parameterization of the household value function (alternatively one could approximate the consumption function).

(2) Solve for the stationary state of the economy without aggregate shocks.
(3) Choose a DSF which uses the information from the solution without aggregate shocks.
(4) Choose a set of statistics m which describe the cross-sectional distribution and which are taken as arguments of the household value function in the model with aggregate shocks.
(5) Solve the model with aggregate shocks by backward induction. Use the stationary distribution as a tool for the DSF.
(6) Optional: simulate the model, and use the distributions arising in the information to obtain a more accurate DSF.
(7) Optional: iterate between 5 and 6 to further improve the DSF, until no increase in accuracy is found.

2.1. Discretization

2.1.1. Representing distributions

In many applications it may be appropriate to use a smooth representation of the cross-sectional density (for example, an exponential of polynomials as in Algan et al., 2008). Here I use a non-smooth approximation, namely a histogram. The reason is that the theoretical model implies a stationary distribution with a large number of discrete point masses, because the idiosyncratic shock has only two realizations. Then the households at the liquidity constraint will enter the range of positive capital as a discrete point mass, if they receive a positive employment shock. In the calibration that we investigate below, the fraction of households at the liquidity constraint is so small that the distribution appears to be almost continuous, but other calibrations imply a distribution with many spikes. My parameterization accounts for this possibility.

A discrete approximation D of a distribution is then given by

(1) a grid $\kappa_D^{\varepsilon,j}$ for $j=0,\ldots,n_d$ and each employment status $\varepsilon \in \{u, e\}$;
(2) a set of probabilities $p_D^{\varepsilon,j}$, $j=1,\ldots,n_d$, $\varepsilon \in \{u, e\}$ with $\sum_{j=1}^{n_d} p_D^{u,j} = u$ and $\sum_{j=1}^{n_d} p_D^{e,j} = 1-u$, which denote the mass of households with employment status ε and capital $k \in (\kappa_D^{\varepsilon,j-1}, \kappa_D^{\varepsilon,j})$.[1]

Underlying this discretization is the assumption that, for each ε, the cross-sectional density is constant within any interval $k \in (\kappa_D^{\varepsilon,j-1}, \kappa_D^{\varepsilon,j})$. The expectation of any function $g(k,\varepsilon)$ over this distribution is therefore given by

$$E_D[g(k,\varepsilon)] \equiv \sum_{\varepsilon\in\{e,u\}} \sum_{j=1}^{n_d} \frac{p_D^{\varepsilon,j}}{\kappa_D^{\varepsilon,j} - \kappa_D^{\varepsilon,j-1}} \int_{\kappa_D^{\varepsilon,j-1}}^{\kappa_D^{\varepsilon,j}} g(k,\varepsilon)\,dk \tag{2}$$

[1] Notice that the $p_t^{\varepsilon,j}$ in Den Haan (2008) are normalized differently, adding up to unity for each ε.