

Solving Heterogeneous-Agent Models by Projection and Perturbation*

Michael Reiter, Universitat Pompeu Fabra, Barcelona

September 2006

Abstract

The paper proposes a numerical solution method for general equilibrium models with a continuum of heterogeneous agents, which combines elements of projection and of perturbation methods. The basic idea is to solve first for the stationary solution of the model, without aggregate shocks but with fully specified idiosyncratic shocks. Afterwards one computes a first-order perturbation of the solution in the aggregate shocks. This approach allows to include a high-dimensional representation of the cross-sectional distribution in the state vector. The method is applied to a model of household saving with uninsurable income risk and liquidity constraints. The model includes not only productivity shocks, but also shocks to redistributive taxation, which cause substantial short-run variation in the cross-sectional distribution of wealth. If those shocks are operative, it is shown that a solution method based on very few statistics of the distribution is not suitable, while the proposed method can solve the model with high accuracy, at least for the case of small aggregate shocks. Techniques are discussed to reduce the dimension of the state space such that higher order perturbations are feasible.

Matlab programs to solve the model can be downloaded.

JEL classification: C63, C68, E21

Keywords: heterogeneous agents; projection methods; perturbation methods

Address of the author:

Michael Reiter
Universitat Pompeu Fabra
Department of Economics and Business
Ramon Trias Fargas, 25–27
08005 Barcelona
Spain

e-mail: michael.reiter@upf.edu

*I am grateful to Ken Judd for very helpful discussions on the topics of this paper. Financial support from Spanish ministry grants SEJ2004-03619 and SEC2002-01601 is gratefully acknowledged.

1 Introduction

Stochastic general equilibrium models with incomplete markets and a continuum of heterogeneous agents pose some hard computational problems. The main problem is that the state vector includes the whole cross-sectional distribution of wealth (or whatever the individual state variable of the agent is). This is an infinite-dimensional object.

The most widely used method to solve this class of models is the one pioneered in Krusell and Smith (1998). Alternative methods that have been around for some years are den Haan (1997) and Reiter (2002). All these approaches have in common that they try to represent the cross-sectional distribution of wealth by a very small number of statistics in order to reduce the dimensionality of the state space. For the models studied in those papers, this approach was shown to work very well. In fact, Krusell and Smith (1998) find that in order to predict future values of aggregate capital with reasonable precision, everything one has to know about today's distribution of capital is the mean. Higher moments of the distribution matter very little. In Krusell and Smith (2006) they explain this approximate aggregation result in detail.

However, this can hardly serve as a general approach to the solution of heterogeneous agent models. For example, in models where firms follow (S,s)-policies (price setting, inventory holdings etc.), we expect that the cross-sectional distribution of the relevant variables enters the solution in an essential way, and the solution algorithm may have to keep track of a medium- or high-dimensional representation of this distribution. Some very recent approaches promise to live up to this challenge. Algan, Allais, and Haan (2006) develop a method that uses a flexible parameterization of the cross-sectional distribution and makes heavy use of projection methods. It increases computational efficiency over earlier approaches in a number of ways and will allow to handle a higher number of state variables. Kim, Kim, and Kollmann (2005) and Preston and Roca (2006) apply perturbation methods. They perturb the solution around the deterministic steady state, with no aggregate or idiosyncratic shocks and a degenerate cross-sectional distribution. The disadvantage of this approach is that it is a local approximation method, and the handling of inequality constraints, for example liquidity constraints, is problematic. The main advantage is that the method will allow to handle higher moments and/or several individual state variables.

In this paper, I propose an alternative procedure, which combines features of projection methods and of perturbation methods. The idea is to compute a solution that is fully nonlinear in the idiosyncratic shocks, but only linear in the aggregate shocks. To do this, we first compute the stationary solution of the model without aggregate shocks.¹ We approximate this solution with a finite parameterization (put into the vector Θ), which describes the individual choice functions and the cross-sectional distribution of wealth. In a second step, we compute an approximation of the dynamics of Θ_t that is linear in the aggregate shocks and in Θ_t itself. With this approach we can include a very detailed representation of the cross-sectional distribution in the state vector. In the numerical examples

¹This part of the solution is similar in spirit to Bohacek and Kejak (2005), although many details of the implementation differ.

of Section 5, the distribution is characterized by up to 1000 state variables. This feature makes the method applicable to a wide range of models, as long as linearity in aggregate shocks is good enough for the purpose at hand. I apply the method to a simple model of household saving with uninsurable income risk and liquidity constraints, to facilitate the exposition of the method and compare it to the earlier literature. The Matlab programs to solve this model are available at “<http://www.econ.upf.es/~reiter/hetagent.tar.gz>”.

Linearity of the solution in aggregate shocks is probably good enough in many cases where individual shocks are much bigger than aggregate shocks, a typical situation in economic applications. But even in cases where nonlinearity in aggregate shocks is essential, the linear solution is a useful first step. With it we can study in detail the dynamics of the cross-sectional distribution that results from small shocks, and we can find out whether some form of approximate state aggregation is possible, and what suitable state variables are. Afterwards we can compute higher-order approximations in this reduced state space.

The plan of the paper is as follows. Section 2 presents a simple model that will serve as a test case for the solution method. The method is presented in detail in Section 3. Section 4 explains how parameters and functional forms were chosen. Numerical results for the simple model are given in Section 5. Section 6 concludes.

2 The Model

To facilitate the exposition of the solution method, I keep the model as simple as possible. There is a continuum of infinitely lived households of unit mass. Households are ex ante identical, and differ ex post through the realization of their individual labor productivity. They supply their labor inelastically. Production takes place in competitive firms with constant-returns-to-scale technology. A government is introduced into this model to the sole purpose of creating some random redistribution of wealth. This helps to identify the effect of the wealth distribution on the dynamics of aggregate capital.

2.1 Production

Output is produced by perfectly competitive firms, using the Cobb-Douglas production function

$$Y_t = \mathcal{Y}(\bar{K}_{t-1}, L_t, t) = Z_t A \bar{K}_{t-1}^\alpha L_t^{1-\alpha}, \quad 0 < \alpha < 1 \quad (1)$$

where A is a constant. Notice that we follow the convention to choose the time index of a variable as the period in which this variable is determined. Production at the beginning of period t uses \bar{K}_{t-1} , the aggregate capital stock determined at the end of period $t - 1$. Since labor supply is exogenous, and individual productivity shocks cancel out due to the law of large numbers, aggregate labor input is constant and normalized to $L_t = 1$.

The aggregate productivity parameter Z_t follows the AR(1) process

$$\log Z_{t+1} = \rho_z \log Z_t + \epsilon_{Z,t+1} \quad (2)$$

The before tax interest rate \bar{r}_t and the wage w_t are determined competitively:

$$\bar{r}_t = \mathcal{Y}_K(\bar{K}_{t-1}, L_t, t) \quad (3)$$

$$w_t = \mathcal{Y}_L(\bar{K}_{t-1}, L_t, t) \quad (4)$$

2.2 The Government

As I said earlier, the only purpose of the government is to create some random redistribution of wealth. The government does this through a tax τ_t levied at the beginning of t on the capital stock accumulated at the end of period $t - 1$. The after tax interest rate r_t is then related to the before tax rate by

$$r_t = \bar{r}_t - \tau_t \quad (5)$$

The revenues are redistributed lump sum, providing a transfer T_t to every household. The government has to balance its budget every period:

$$T_t = \tau_t \bar{K}_{t-1} \quad (6)$$

The tax rate follows an AR(1) process around its steady state value τ^* :

$$\tau_{t+1} - \tau^* = \rho_\tau(\tau_t - \tau^*) + \epsilon_{\tau,t+1} \quad (7)$$

2.3 The Household

There is a continuum of households, indexed by i . Households differ ex post by their labor productivity $\xi_{t,i}$. Labor productivity is assumed to be i.i.d., and is normalized to have unit mean:

$$\mathbb{E}_{t-1} \xi_{t,i} = 1 \quad (8)$$

For convenience of exposition I assume that $\xi_{t,i}$ has a continuous distribution with density function $f(\xi)$. The corresponding distribution function is denoted by $F(\xi)$. The density and distribution function of labor income depend on the wage and are denoted by

$$f_t^y(y) = \frac{1}{w_t} f\left(\frac{y}{w_t}\right) \quad (9a)$$

$$F_t^y(y) = F\left(\frac{y}{w_t}\right) \quad (9b)$$

The household supplies inelastically one unit of labor. Its labor earnings are therefore given by

$$\bar{y}_{t,i} = w_t \xi_{t,i} \quad (10)$$

With the government lump sum transfer T_t , the net non-capital income is given by

$$y_{t,i} = \bar{y}_{t,i} + T_t \quad (11)$$

Household i enters period t with asset holdings $k_{t-1,i}$ left at the end of the last period. It receives the net interest rate r_t on its assets, such that the available resources after income of period t (“cash on hand”) are given by

$$x_{t,i} = (1 + r_t)k_{t-1,i} + y_{t,i} \quad (12)$$

$$(13)$$

Cash on hand is split between consumption and asset holdings:

$$k_{t,i} = x_{t,i} - c_{t,i} \quad (14)$$

We impose the liquidity constraint

$$k_{t,i} \geq \underline{k} \quad (15)$$

The solution of the household problem is given by a consumption function $C(x; \Omega_t)$, or equivalently a savings function

$$K_t(x) \equiv x - C(x; \Omega_t) \quad (16)$$

where Ω_t denotes the vector of state variables, which will be specified later. The first order equation of the household problem is the Euler equation

$$\begin{aligned} U'(C(x; \Omega_t)) &\geq \beta \mathbf{E}_t [R(\Omega_{t+1})U'(C(R(\Omega_{t+1})(x - c_{t,i}) + y_{t+1,i}; \Omega_{t+1})))] \\ \text{and } C(x; \Omega_t) &= x - \underline{k} \end{aligned} \quad (17a)$$

or

$$\begin{aligned} U'(C(x; \Omega_t)) &= \beta \mathbf{E}_t [R(\Omega_{t+1})U'(C(R(\Omega_{t+1})(x - c_{t,i}) + y_{t+1,i}; \Omega_{t+1})))] \\ \text{and } C(x; \Omega_t) &< x - \underline{k} \end{aligned} \quad (17b)$$

where $R(\Omega_{t+1}) \equiv 1 + r(\Omega_{t+1})$. The expectation in (17) is over the distribution of $\xi_{t+1,i}$ and of Ω_{t+1} . Since the household problem is concave, Equ. (17) together with the constraint (15) are both necessary and sufficient for a solution of the household problem. Notice that the solution depends on the aggregate state Ω_t only through the expected future dynamics of the factor prices.

2.4 Dynamics of the Cross-Sectional Distribution

The cross-sectional distribution of cash on hand at the beginning of period t , after the realization of labor income, is continuous, because we have assumed that idiosyncratic productivity $\xi_{t,i}$ has a continuous distribution. Denote the density of cash-on-hand by $\phi_t(\xi)$. The distribution of assets at the end of period t has a mass point of liquidity constrained households which hold the minimum level of assets \underline{k} . For $k > \underline{k}$, the distribution is continuous and we denote the density function by $\psi_t(k)$. The cross-sectional distribution function $\Psi_t(\tilde{k})$ is defined as the fraction of households i such that $k_{t,i} \leq \tilde{k}$. The mass point satisfies

$$\Psi_t(\underline{k}) = \int_{-\infty}^{\chi^t} \phi_t(x) \, dx \quad (18a)$$

Recall that χ_t is the value of cash-on-hand where the liquidity constraint starts binding. The cross-sectional density of capital $\psi_t(i)$ satisfies

$$\psi_t(K_t(x)) K_t'(x) = \phi_t(x), \quad x > \chi_t \quad (18b)$$

For the following formula, we use the abbreviation

$$a_t(k) \equiv (1 + r_t)k + T_t \quad (19)$$

for the level of assets after interest and lump sum subsidy, but before labor income. Then the density of cash-on-hand is related to last period's distribution of capital by

$$\phi_t(x) = \Psi_{t-1}(\underline{k}) f_t^y(x - a_t(\underline{k})) + \int_{\underline{k}}^{\infty} f_t^y(x - a_t(k)) \psi_{t-1}(k) \mathbf{d}k \quad (20)$$

Equs. (18) and (20) can be combined into a dynamic equation in the distribution or of k .² Using

$$\begin{aligned} \int_{-\infty}^{\chi_t} \int_{\underline{k}}^{\infty} f_t^y(x - a_t(k)) \psi_{t-1}(k) \mathbf{d}k \mathbf{d}x &= \int_{\underline{k}}^{\infty} \psi_{t-1}(k) \int_{-\infty}^{\chi_t} f_t^y(x - a_t(k)) \mathbf{d}x \mathbf{d}k \\ &= \int_{\underline{k}}^{\infty} \psi_{t-1}(k) F_t^y(\chi_t - a_t(k)) \mathbf{d}k \end{aligned} \quad (21)$$

we get

$$\Psi_t(\underline{k}) = \Psi_{t-1}(\underline{k}) F_t^y(\chi_t - a_t(\underline{k})) + \int_{\underline{k}}^{\infty} \psi_{t-1}(k) F_t^y(\chi_t - a_t(k)) \mathbf{d}k \quad (22a)$$

and

$$\psi_t(K_t(x)) K_t'(x) = \Psi_{t-1}(\underline{k}) f_t^y(x - a_t(\underline{k})) + \int_{\underline{k}}^{\infty} f_t^y(x - a_t(k)) \psi_{t-1}(k) \mathbf{d}k, \quad \forall x > \chi_{t-1} \quad (22b)$$

2.5 Equilibrium of the Model

We call the set of equations (2), (6), (7), (17) and (22) the **Theoretical Model**. The “natural” state variables of this model at time t are the exogenous driving forces Z_t and τ_t and the cross-sectional distribution of capital holdings inherited from the last period, $\Psi_{t-1}(k)$. We therefore choose $\Omega_t \equiv (Z_t, \tau_t, \Psi_{t-1}(k))$. A stationary equilibrium in Ω_t consists of a consumption function $C(x; \Omega_t)$, a stochastic process of cross sectional distribution functions $\Psi_t(k)$ and a process of lump sum transfers T_t such that

1. the consumption function satisfies the Euler equation (17)

²We could also write the dynamic equation in the distribution of x , but it is better to write it in k , because the distribution of k at the end of period $t - 1$ is the appropriate state variable in period t . The distribution of x in t is already affected by the aggregate shocks of period t .

2. the cross sectional distribution satisfies the dynamic equations (22)
3. the transfer satisfies the government budget constraint (6).

It has not yet been proven that a stationary equilibrium in Ω_t exists; Miao (2006) proves a recursive equilibrium in a larger state space including expected payoffs to households. In the following section, we will introduce a discrete approximation to this economy, and the computations can probably be seen as a “constructive proof” of the equilibrium of the approximate economy for sufficiently small fluctuations of the aggregate variables Z_t and τ_t (to make this more formal, one would have to do a theoretical analysis of the perturbation approach along the lines of Jin and Judd (2002) or Judd (2004), but this is not the focus of this paper).

3 Method

The solution method proposed here has three steps:

1. Providing a finite representation of the economy, in particular of the consumption function $C(x; \Omega_t)$ and the cross sectional distribution of wealth $\Psi_t(k)$.
2. Computing the steady state of the economy, that is, the stationary economy when the aggregate shocks are identically zero. Concretely, aggregate productivity Z_t and the tax on capital τ_t are constant. The idiosyncratic productivity shocks $\xi_{t,i}$ are the same as in the full model.
3. Computing a first-order perturbation of the steady state solution in the aggregate variables Z_t and τ_t . This gives a precise solution of the finite-dimensional model for small aggregate shocks $\epsilon_Z, \epsilon_\tau$.

3.1 Finite Approximation of the Model Equations

3.1.1 Savings, Consumption and the Euler Equation

The consumption function $C(x; \Omega_t)$ is very well behaved. At least for the steady state case, we know (Carroll 2004) that

1. $C(x; \Omega_t)$ is concave in x .
2. There is an χ_t such that $C(x; \Omega_t) = x - \underline{k}$ for $x \leq \chi_t$ and $C(x; \Omega_t) < x - \underline{k}$ for $x > \chi_t$
3. The limit $\lim_{x \rightarrow \infty} \frac{\partial C(x; \Omega_t)}{\partial x}$ exists.

This motivates the following smooth approximation of the consumption or the savings function. I represent the savings function $K_t(x)$ by $n_p + 1$ numbers, collected into the vector \mathbf{c}_t : the critical value χ_t , and the values $k_{t,i}$, $i = 1, \dots, n_p$ at n_p knot points $x_{t,i}$ with $x_{t,i} > \chi_t$. The knot points are chosen as $x_{t,0} = \chi_t$ and $x_{t,i} = \chi_t + X_i$, $i = 1, \dots, n_p$

with some fixed set of grid points X_i .³ Off the knot points $x_{t,i}$, the function $K_t(x)$ is then approximated by

$$\hat{K}(x; \mathbf{c}_t) \approx \begin{cases} \underline{k} & \text{for } x \leq \chi_t \\ CSI(x) & \text{for } \chi_t < x \leq x_{t,n_p} \\ k_{t,n_p} + CSI'(x_{t,n_p})(x - x_{t,n_p}) & \text{for } x > x_{t,n_p} \end{cases} \quad (23)$$

where $CSI(x)$ stands for ‘‘cubic spline interpolation’’, the natural cubic spline that interpolates the points (χ_t, \underline{k}) , $(x_{t,i}, k_{t,i})$, $i = 1, \dots, n_p$. Beyond the last knot point x_{t,n_p} , we approximate the savings function by a straight line, with the slope given by the derivative of the spline at k_{t,n_p} .

We approximate the expectation over the individual productivity shock ξ on the rhs of the Euler equation (17) by a finite sum, with quadrature points $\hat{\xi}_i$ and weights ω_i^ξ , $i = 1, \dots, n_\xi$. The details of the quadrature depend on the distribution of ξ , and will be discussed in Section 4.2.

Since we approximate the saving function by a spline with $n_p + 1$ degrees of freedom, we apply a collocation method and require the Euler equation (17) to hold with equality at the knot points $x_{t,i}$:

$$U'(\hat{C}(x_{t-1,i}; \mathbf{c}_{t-1})) = \beta \sum_{j=1}^{n_\xi} \omega_j^\xi \left[\hat{R}(\mathbf{p}_t) U'(\hat{C}(\hat{X}_{ij}; \mathbf{c}_t)) \right] + \eta_{i,t}^c, \quad i = 0, \dots, n_p \quad (24a)$$

where $\hat{C}(x; \mathbf{c}_t)$ is short for $x - \hat{K}(x; \mathbf{c}_t)$ and

$$\hat{X}_{ij} \equiv \hat{R}(\mathbf{p}_t) \left(x_{t-1,i} - \hat{C}(x_{t-1,i}; \mathbf{c}_{t-1}) \right) + \hat{W}(\mathbf{p}_t) \hat{\xi}_j + T_t \quad (24b)$$

Equ. (24a) uses the notation of Sims (2001): the $\eta_{i,t}^c$ are expectation errors that are determined endogenously in the solution of the system (cf. Section 3.3).

3.1.2 Wealth Distribution

There are several ways to parameterize a distribution function. One approach is to approximate the density function by a linear combination of known, smooth basis functions (polynomials, Gaussian densities, etc.) and force the dynamic equations (22) to hold on a set of collocation points. This is the approach taken in Bohacek and Kejak (2005). The downside of this approach is that we have to take care to enforce the non-negativity of the density function, which will involve the solution of a set of nonlinear equations. Here I follow a different path, which allows to compute the stationary distribution by a system of linear equations (cf. (33)). I first truncate the distribution of capital at a maximum level \bar{k} . Then I split the support into n_d small intervals by a grid of points $\underline{k} = \kappa_0, \kappa_1, \dots, \kappa_{n_d} = \bar{k}$

³For the X_i I take a grid with many points close to 0, because the consumption function has high curvature close to the kink point χ_t and is difficult to approximate with splines there.

(I will choose $n_d = 1000$ below). The truncation points \bar{k} must be chosen such that in the steady state distribution, a very small fraction of households is close to \bar{k} . Then I characterize the distribution by the mass at the lower bound, $p_t^0 \equiv \Psi_t(\underline{k})$, and the mass on the intervals $p_t^i \equiv \Psi_t(\kappa_i) - \Psi_t(\kappa_{i-1})$, $i = 1, \dots, n_d$, which are stacked into the vector $\mathbf{p}_t \equiv (p_t^0, p_t^1, \dots, p_t^{n_d})$. These probabilities satisfy

$$\begin{aligned}
p_t^i &= \int_{\kappa_{i-1}}^{\kappa_i} \psi_t(k) \, \mathbf{d}k = \int_{x_{t,i-1}}^{x_{t,i}} \psi_t(K_t(x)) K_t'(x) \, \mathbf{d}x \\
&= \int_{x_{t,i-1}}^{x_{t,i}} \left[\Psi_{t-1}(\underline{k}) f_t^y(x - a_t(\underline{k})) + \int_{\underline{k}}^{\infty} f_t^y(x - a_t(k)) \psi_{t-1}(k) \, \mathbf{d}k \right] \mathbf{d}x \\
&= \Psi_{t-1}(\underline{k}) [F_t^y(x_{t,i} - a_t(\underline{k})) - F_t^y(x_{t,i-1} - a_t(\underline{k}))] \\
&\quad + \int_{\underline{k}}^{\infty} \psi_{t-1}(k) [F_t^y(x_{t,i} - a_t(k)) - F_t^y(x_{t,i-1} - a_t(k))] \, \mathbf{d}k
\end{aligned} \tag{25}$$

where the $x_{t,i}$ are implicitly defined by

$$K_t(x_{t,i}) = \kappa_i \tag{26}$$

Formula (25) contains expressions of the form $\int_{\underline{k}}^{\infty} \psi_{t-1}(k) F_t^y(x - a_t(k)) \, \mathbf{d}k$, which have to be evaluated for many different x . To compute this integral, I will assume that the probability mass is approximately uniformly distributed within each interval (constant density function):

$$\psi_t(k) \approx \frac{p_t^i}{\kappa_i - \kappa_{i-1}}, \quad k \in (\kappa_{i-1}, \kappa_i), \forall i \tag{27}$$

Then we get, using (19), that

$$\begin{aligned}
\int_{\underline{k}}^{\infty} \psi_{t-1}(k) F_t^y(x - a_t(k)) \, \mathbf{d}k &= \sum_{i=1}^{n_d} \int_{\kappa_{i-1}}^{\kappa_i} \psi_{t-1}(k) F\left(\frac{x - a_t(k)}{w_t}\right) \, \mathbf{d}k \\
&\approx \sum_i \frac{p_{t-1}^i}{\kappa_i - \kappa_{i-1}} \int_{\kappa_{i-1}}^{\kappa_i} F\left(\frac{x - a_t(k)}{w_t}\right) \, \mathbf{d}k \\
&= \sum_i \frac{p_{t-1}^i}{\kappa_i - \kappa_{i-1}} \int_{(x - a_t(\kappa_i))/w_t}^{(x - a_t(\kappa_{i-1}))/w_t} F(\xi) \frac{w_t}{R_t} \, \mathbf{d}\xi \\
&= \sum_i \frac{p_{t-1}^i}{\kappa_i - \kappa_{i-1}} \frac{w_t}{R_t} \left[F^{(1)}\left(\frac{x - a_t(\kappa_{i-1})}{w_t}\right) - F^{(1)}\left(\frac{x - a_t(\kappa_i)}{w_t}\right) \right]
\end{aligned} \tag{28}$$

Here, $F^{(1)}(x)$ is defined as the integral of the distribution function $F(x)$ with respect to x . To evaluate it quickly, we do the following preparatory steps at the beginning of the model solution. First, we compute $F^{(1)}(x)$ at a dense grid of x by numerical integration. Then we use those values to compute a spline approximation of $F^{(1)}(x)$ which preserves the convexity of this function. I use Schumaker's quadratic splines (Judd 1998, p.231),

exploiting that the derivative of $F^{(1)}(x)$ is $F(x)$. Afterwards we evaluate $F^{(1)}(x)$ always by spline interpolation.

These derivations are summarized in the following linear dynamic equation for \mathbf{p}_t :

$$\mathbf{p}_t = \Pi(\Omega_t)\mathbf{p}_{t-1} \quad (29a)$$

where the elements of the transition matrix $\Pi(\Omega_t)$ are (cf. (22a), (25) and (28))

$$\Pi_{0,0}(\Omega_t) = F_t^y(\chi_t - a_t(\underline{k})) \quad (29b)$$

$$\Pi_{i,0}(\Omega_t) = F_t^y(x_{t,i} - a_t(\underline{k})) - F_t^y(x_{t,i-1} - a_t(\underline{k})), \quad i = 1, \dots, n_d \quad (29c)$$

$$\Pi_{0,j}(\Omega_t) = \frac{1}{\kappa_j - \kappa_{j-1}} \frac{w_t}{R_t} \left[F^{(1)}\left(\frac{\chi_t - a_t(\kappa_{j-1})}{w_t}\right) - F^{(1)}\left(\frac{\chi_t - a_t(\kappa_j)}{w_t}\right) \right], \quad j = 1, \dots, n_d \quad (29d)$$

$$\begin{aligned} \Pi_{i,j}(\Omega_t) = \frac{1}{\kappa_j - \kappa_{j-1}} \frac{w_t}{R_t} & \left[F^{(1)}\left(\frac{x_{t,i} - a_t(\kappa_{j-1})}{w_t}\right) - F^{(1)}\left(\frac{x_{t,i} - a_t(\kappa_j)}{w_t}\right) \right. \\ & \left. - F^{(1)}\left(\frac{x_{t,i-1} - a_t(\kappa_{j-1})}{w_t}\right) + F^{(1)}\left(\frac{x_{t,i-1} - a_t(\kappa_j)}{w_t}\right) \right], \quad i, j = 1, \dots, n_d \end{aligned} \quad (29e)$$

Notice that $\Pi(\Omega_t)$ depends on $\hat{K}(x; \mathbf{c}_t)$ through (26).

3.1.3 The Discrete Model

With the finite approximations defined above, we can now reduce the model to a set of finite equations in each period t . The **Discrete Model** consists of the equations (2), (7), (24), (29) and the government budget constraint

$$T_t = \tau_t \hat{K}(\mathbf{p}_t) \quad (30)$$

where $\hat{K}(\mathbf{p}_t)$ is aggregate capital, computed using the following formula for the j -th moment of the capital stock

$$\int k^j d\Psi_t(k) = \underline{k}^j \Psi_t(\underline{k}) + \int_{\underline{k}}^{\infty} k^j \psi_t(k) \mathbf{d}k \approx \underline{k}^j p_t^0 + \sum_i p_t^i \frac{\kappa_i^{j+1} - \kappa_{i-1}^{j+1}}{(j+1)(\kappa_i - \kappa_{i-1})} \quad (31)$$

Equ. (31) uses again the approximation (27).

These equations define, for each period t , a system of $n_p + n_d + 4$ equations in just as many variables: \mathbf{c}_t , \mathbf{p}_t (note that p_t^0 is redundant since probabilities sum to unity), T_t , Z_t and τ_t . Of those, \mathbf{p}_t are endogenous state variables, Z_t and τ_t are exogenous state variables, while \mathbf{c}_t and T_t are control or “jump” variables.

3.2 Solving For the Steady State

Given the finite approximation of the model described above, we are now ready to solve for the stationary state of the model, characterized by constant values of the exogenous aggregate variables $Z_t = 1$ and $\tau_t = \tau^*$.

Given a tax rate τ^* , solving for the steady state involves the following steps:

1. For any guess of the steady state capital stock K^* , take the net interest rate $r^* = \mathcal{Y}_K(K^*, 1, 1) - \tau^*$, wage $w^* = \mathcal{Y}_L(K^*, 1, 1)$ and transfers $T^* = \tau^* K^*$. Then compute the parameters of the savings and consumption function, \mathbf{c}^* , so as to satisfy the household optimality condition (24).
2. For any r^* , T^* and \mathbf{c}^* , compute the parameters of the cross-sectional distribution \mathbf{p}^* so as to satisfy (29).
3. Find the capital level K^* that is consistent with the cross-sectional distribution of capital computed in the last step.

These steps are now described in more detail.

3.2.1 Finding the Parameters of the Consumption Function

In steady state, the Euler equation (24) reduces to

$$U' \left(\hat{C}(x_i^*; \mathbf{c}^*) \right) = \beta \sum_{j=1}^{n_\xi} \omega_j^\xi \left[(1 + r^*) U' \left(\hat{C}((1 + r^*) k^*(x_i^*) + w^* \hat{\xi}_j + T^*; \mathbf{c}^*) \right) \right],$$

$$i = 0, \dots, n_p \quad (32)$$

Given r^* , w^* and T^* , this is a system of $n_p + 1$ nonlinear equations in the $n_p + 1$ components of \mathbf{c}^* . This can be solved without problems by a quasi-Newton algorithm (cf. the routine 'broydn.m' in the accompanying Matlab code).

3.2.2 Computing the Invariant Distribution

Given the steady state savings function $\hat{K}(x; \mathbf{c}^*)$, we obtain from (29) the matrix Π^* that describes the stationary transition dynamics of the wealth distribution. The finite parameterization of the steady state distribution \mathbf{p}^* is then a nonzero solution of the linear equation system

$$\mathbf{p}^* = \Pi^* \mathbf{p}^* \quad (33)$$

From the theory of Markov chains, and the properties of the consumption function, one can show that there is a unique \mathbf{p}^* such that $\sum_i \mathbf{p}_i^* = 1$.⁴ This means that the null space of $(\Pi^* - I)$ has dimension one. In this case, a fast way to compute an element of this null

⁴Even the richest household ($k = \bar{k}$) is driven down to $k = \underline{k}$ in a finite number of period if it always draws the worst possible income shock. Starting from \underline{k} , the household reaches \bar{k} in a few periods, if it

space is by “inverse iteration”, using LU-decomposition (Press, Flannery, Teukolsky, and Vetterling 1986, Section 11.7).⁵ The non-negativity of \mathbf{p}^* is assured since Π is a Markov transition matrix.

3.2.3 Solving for the Aggregate Capital Stock

Starting from a guess of the aggregate capital stock K^* , we have found \mathbf{p}^* , which implies a value for K^* by (31). Therefore we are left with a one-dimensional root-finding problem in K^* . This can be easily solved by bisection or, faster, by Brent’s method (cf. Press, Flannery, Teukolsky, and Vetterling (1986, Section 9.3), which is implemented in the Matlab command ‘fzero’).

3.3 Perturbation of the Steady State

In Section 3.1.3 we saw that Equations (2), (7), (24), (29) and (30) define a system of $n_p + n_d + 4$ equations in just as many variables. Collect all those variables in the vector Θ_t . Collect the expectation errors in the vector η_t , and the exogenous shocks in $\epsilon_t \equiv (\epsilon_{Z,t}, \epsilon_{\tau,t})$. The system of equations can then be written in the compact form

$$H(\Theta_{t-1}, \Theta_t, \eta_t, \epsilon_t) = 0 \quad (35)$$

Denote by Θ^* the steady state values of Θ_t . Obviously, we have

$$H(\Theta^*, \Theta^*, 0, 0) = 0 \quad (36)$$

The linearized stochastic model is given by

$$H_1(\Theta^*, \Theta^*, 0, 0) (\Theta_{t-1} - \Theta^*) + H_2(\Theta^*, \Theta^*, 0, 0) (\Theta_t - \Theta^*) + H_3 \eta_t + H_4 \epsilon_t = 0 \quad (37)$$

where H_i , $i = 1, 2, 3, 4$ denotes the partial derivative of H with respect to its i -th argument. The analytical computation of the partial derivatives would be tedious. Automatic differentiation would be possible, but is not available in the Matlab environment that I use. I

always draws the highest income shocks. Since the income distribution has a connected support, and the consumption function is continuous, the household reaches all intermediate capital values with positive probability in the same number of periods. Then Π^* satisfies all the conditions of Stokey and Lucas (1989, Theorem 11.4).

⁵Briefly, it works as follows. Choose randomly an n_d -vector b and solve

$$(\Pi - I)x = b \quad (34)$$

for x by use of an LU-decomposition. The operator $(\Pi - I)$ is singular, and b lies “with probability 1” not in its range. When trying to solve (34), the LU-decomposition will encounter a zero or almost-zero pivot, which it replaces by some very small number (not all implementations of the LU-decomposition do this, but the Matlab ‘\’ operator seems to do it). It then returns an x that lies, almost to machine precision, in the null space. One can check whether x satisfies $(\Pi - I)x = 0$ with sufficient precision, and add a correction step if necessary. In my experience so far, this was never necessary. Then set $\mathbf{p}^* = x / \sum_i x_i$. In the numerical experiments reported below, I have done this for matrices up to $n_d = 5000$. In cases where n_d is much bigger (and hopefully sparse!), one should use more sophisticated procedures to find the invariant distribution (see for example Virnik (2006)).

therefore choose the simplest approach, computing the derivatives by forward differencing. This is clearly not computationally efficient compared to automatic differentiation, but it turns out to be good enough.

The system of equations (37) is in the form that can be solved by the package of Sims (2001). The number of state variables is $n_d + 2$. In the numerical examples below, I use $n_d = 1000$. Sims' package then needs about 30 minutes to solve the model (most of the work is in the reordering of eigenvectors, which could perhaps be implemented more efficiently).

3.4 State Space Reduction

3.4.1 General Description

With our non-parametric approach, to solve precisely for the steady state distribution we need a fine grid, that is a big n_d . In the model with aggregate shocks, the n_d parameters of the cross-sectional distribution, contained in \mathbf{p}_t , are then all state variables. For very high n_d , say 10000, one can still solve for the steady state \mathbf{p}^* , but solving for the dynamics of the distribution is currently not feasible on a PC. Even if feasible, one might get a sufficiently precise solution by some form of state space reduction: approximate the perturbations in n_d -dimensional space by a lower-dimensional smooth parameterization, for example by the linear combination of some smooth basis functions:

$$\mathbf{p}_t \approx \mathbf{p}^* + B\mathbf{b}_t \quad (38)$$

Here B is a $n_d \times n_B$ -matrix of known basis functions, and \mathbf{b}_t is the vector of time-varying coefficients. In (38), we separate the parameterization of the invariant distribution \mathbf{p}^* from the parameterization of the deviations, $B\mathbf{b}_t$. The vector of state variables characterizing the deviations of the distribution from the steady state is now \mathbf{b}_t with dimension n_B .

The matrix B represents a linear operation that transforms an n_B -dimensional parameter vector into an n_d -dimensional distribution. If $n_B < n_d$, we cannot expect (29a) to hold for all n_d points of \mathbf{p}_t . To define exactly n_B conditions on the distribution, we define an inverse operation to B , which assigns to any n_d -dimensional distribution an n_B -dimensional parameter vector. Denote the matrix of this operation by P_B . Then we can replace (29a) by the weaker condition

$$\mathbf{b}_{t+1} = P_B [\Pi(\Omega_t) (\mathbf{p}^* + B\mathbf{b}_t) - \mathbf{p}^*] \quad (39)$$

Obviously, we will require that $P_B B\mathbf{b}_t = \mathbf{b}_t$. This means that P_B gives the coefficients of a projection into $\text{span}(B)$. One could choose the orthogonal projection, $P_B = (B'B)^{-1}B$. This is not optimal in general: the numerical experiments show that it is important to guarantee that this operation preserves the mean of the distribution. More generally, we may want to impose that the projection preserves a set of linear conditions:

$$M'\mathbf{p} = M'BP_B\mathbf{p}, \quad \forall \mathbf{p} \in \mathfrak{R}^{n_d+1} \quad (40)$$

for some given matrix M . This can be achieved by including M into B , or by choosing

$$P_B = \begin{pmatrix} N'B'B \\ M'B \end{pmatrix}^{-1} \begin{pmatrix} N'B' \\ M' \end{pmatrix}, \quad N \equiv \text{null}(M'B) \quad (41)$$

I have done some experimentation with this, and it seemed best to use (41) and to impose only the preservation of the first moment, not of higher moments.

To summarize, the **Reduced Model** consists of the equations (2), (7), (24), (30) and (39).

3.4.2 Choice of Basis Functions

For the basis functions (the columns of B), it is natural to first try some standard smooth functions, such as a polynomial basis or, what I am going to use below, a spline basis.⁶ We will see in Section 5 that this is a successful strategy.

If we are satisfied with a first-order perturbation in the aggregate state variables, as we are in this paper, we can stop at this point. In other applications it may be necessary to investigate the nonlinearity of the solution in the aggregate shocks. In this case, the first-order perturbation may serve as a useful starting point. Once we have obtained a solution for some basis B , we can try and see whether it is possible to reduce the number of state variables even further without sacrificing much in accuracy. If this turns out to be the case, one can afterwards try a higher-order perturbation in those state variables.

There are probably many ways to pursue this. In the following, I will test only one simple approach:

1. Solve the model for some spline basis B with dimension n_B .
2. Simulate the model for many periods.
3. Analyze the parameter vectors \mathbf{b}_t of the cross-sectional distributions that arise in this simulation, and find out whether they are approximately contained in some subspace with dimension smaller than n_B . This can be done, for example, by a principal components analysis of the set of \mathbf{b}_t . Choose the first $n_b < n_B$ principal components, and stack them into the matrix B_{PC} .
4. Choose the n_b -dimensional basis $B \cdot B_{PC}$ as the new basis, and solve the model again.

4 Parameter Values and Functional Forms

I solve the model at annual frequency. For the model parameters I use standard values, $\beta = 0.95$, $\alpha = 1/3$, $\delta = 0.1$. For the utility function I use CRRA

$$U(c) = \frac{c^{1-\gamma} - 1}{1-\gamma} \quad (42)$$

⁶More precisely, since capital extends over a large state space, I use splines not in the original k , but in the nonlinear transformation $g = \log(k + 1)$.

with risk aversion parameter $\gamma \in \{1, 5\}$. The parameter A in the production function is set to $A = (\beta^{-1} - 1 + \delta)/\alpha$, such that the steady state capital stock of the representative agent model equals 1. This allows us to see the effect of non-insurable idiosyncratic shocks on the capital accumulation in steady state.

For the technology shock I choose $\rho_z = 0.8$ and $\sigma_{\rho_z} = 0.014$, which are standard values, in annual terms. I choose the tax shock as uncorrelated, $\rho_\tau = 0$, to create unpredictable short-run redistributions. The variability of the tax shock is set, rather arbitrarily, to $\sigma_{\rho_\tau} = 0.02$. Taxes fluctuate around zero, so $\tau^* = 0$. Note that the variance of the shocks has no other role here than to scale the simulations, since our approximate solution is linear in the aggregate shocks.

4.1 The Distribution of Individual Productivity

I look at two specifications of the labor income distribution. The first one is the log-normal distribution, which has a thin tail. Following Carroll (2001), I choose standard deviation and mean of $\log(\xi_{t,i})$ as $\sigma = 0.2$ and $\mu = -0.5\sigma^2$, which gives an expected value of productivity of 1, as required by (8).

The second distribution is the Pareto distribution, which has a thick tail. The distribution function is given by

$$F(\xi) = 1 - C\xi^{-m}, \quad m > 1, \quad (43)$$

The distribution has support $\xi \in (C^{1/m}, \infty)$. Choosing the constant $C = ((m-1)/m)^m$ gives an expected value of 1. I choose $m = 2$, which is approximately what Saez (2001) finds for the US income distribution. For $m \leq 2$, the Pareto distribution has infinite variance. In the computations, however, I truncate the distribution at the point $\bar{\xi}$ where $F(\bar{\xi}) = 1 - 10^{-5}$, so that existence of moments is not an issue.

Comparing results for these two distributions with very different characteristics allows us to study the role of the wealth distribution in the dynamics of aggregate capital.

4.2 Quadrature Formulas

In the case of a log-normal distribution, I use Gauss-Legendre quadrature for the $\hat{\xi}_i$ and ω_i^ξ in Equ. (24a). This is implemented in Matlab by the routine 'qwnorm' of the Compecon toolkit that comes with Miranda and Fackler (2002).

In the case of a Pareto distribution, I use the following procedure: I choose a set of cumulative probabilities P_j , $j = 0, \dots, n_\xi$ with $P_0 = 0$, $P_j > P_{j-1}$ and $P_{n_\xi} = 1 - 10^{-5}$ (recall that I truncate the distribution at $1 - 10^{-5}$). Then I set

$$\omega_j^\xi \equiv P_j - P_{j-1} \quad (44a)$$

$$\hat{\xi}_j \equiv \frac{1}{P_j - P_{j-1}} \int_{q(P_{j-1})}^{q(P_j)} \xi f(\xi) \, d\xi \quad (44b)$$

where $q(p)$ is the p -th percentile of the distribution of ξ . The quadrature points $\hat{\xi}_j$ are the expectations of ξ , conditional on $\xi \in (q(P_{j-1}), q(P_j))$. I choose the P_j such that the ω_j^ξ are 10^{-5} , 10^{-4} , 10^{-3} and 10^{-2} , respectively, for the highest intervals. The rest of the probability mass is divided into equal intervals. This makes sure that the top income earners are well represented in the discrete approximation. The numerical experiments reported in Section 5 use $n_d = 41$. Increasing this parameter has no significant effect on the results.

Note that this quadrature rule is only used for the Euler equations. The dynamics of the wealth distribution, developed in Sections 2.4 and 3.1.2, uses the distribution function $F(\xi)$.

5 Numerical Examples

The numerical examples presented here illustrate how the method works with the simple model of Section 2, and which level of accuracy we can expect. After showing some of the dynamic properties of the model, I will investigate the accuracy of the computed solution, proceeding in two steps. First, I estimate the approximation error that comes from replacing the theoretical model of Section 2.5 by the discrete model of Section 3.1.3. Then I estimate the approximation error that arises from the state space reduction (Section 3.4). The error analysis should be understood as applying to the case of (infinitesimally) small aggregate shocks, even if the shocks are scaled to realistic size for convenience. I do not investigate the error arising from big aggregate shocks.

5.1 The Dynamics of the Economic Model

Table 1 shows some summary statistics (mean, standard deviation, skewness and kurtosis) of the cross-sectional distribution of income and wealth. The lognormal distribution has a standard deviation of about 0.2. It is skewed to the right, and has thin tails (the kurtosis is only slightly above that of the normal distribution, which equals 3). In contrast, the truncated Pareto distribution has much higher variance, is strongly skewed to the right and has fat tails (high kurtosis). Notice that, without truncation, these moments would be infinite. The wealth accumulation process levels out some of these differences: with the lognormal, the wealth distribution has higher skewness and kurtosis than the income distribution, the opposite happens with the Pareto distribution. The last column in the table shows the fraction of households that are close to the liquidity constraint ($k \leq \underline{k} + 0.1$). We see first that this fraction is smaller when consumers are more risk averse ($\gamma = 5$), since these consumers try harder to avoid the region of low asset holdings. More interestingly, this fraction is much higher when the income distribution is very unequal and has fat tails (Pareto). Households with very high transitory income shocks save approximately the fraction $1 - r$ of this income. With fat tails, there are more of those people every period. This reduces the interest rate and drives more households down to the liquidity constraint.

Table 2 and Figure 1 illustrate the propagation mechanisms of the model. For this

purpose, both types of shocks were assumed i.i.d. ($\rho_Z = \rho_\tau = 0$). As one would expect, the technology shock creates little variation in the distribution (variance), but a lot of variation in the mean of capital. The tax shock does the opposite. Interestingly, in the case of a distribution with fat tails, the variability of the cross-sectional variance is not as high as in the thin-tail distribution, but it generates a higher variability of mean capital. This is because there are more households in the region where the consumption function has strong curvature (cf. last paragraph), and therefore the wealth distribution matters more for the dynamics of the capital stock.

Figure 1 measures the persistence of fluctuations in capital. The solid lines show the autocorrelation of the capital stock if only the technology shock is active, the dashed ones refer to the case where only the tax shock is active. We see that i.i.d. redistributive shocks (tax shocks) generate fluctuations of capital that are highly persistent, if small.

5.2 The Approximation Error From Discretization

For the discretization, I characterize the distribution by a vector of 1000 probabilities ($n_d = 1000$) and the spline for the consumption function has 100 knot points ($n_p = 100$). Table 3 shows some statistics to estimate the error that we make by this discretization. The first two columns document the residual of the household Euler equation (24a). It is computed as $\frac{1}{C_t} [C_t - U'^{-1}(\beta(1+r^*)E_t(U'(C_{t+1})))]$, such that it expresses the residual as the relative error in consumption. The average absolute residual is computed as the unweighted average on the range between χ^* and $10\bar{K} + y^{max}$, ten times average capital plus the maximum realization of income. We see that the average residual is indeed very low, but the maximum residual is about 3 orders of magnitude higher. The maximum occurs in the region near the critical point where the liquidity constraint starts binding. Here the consumption function has high curvature and is more difficult to approximate by a spline, although I have allocated half of all the knot points of the spline in this region. One can further reduce the residual by increasing the number of parameters, n_ξ , but the accuracy achieved here seems good enough.

The discretization of the cross-sectional distribution of wealth developed in Section 3.1.2 implies two types of errors. The first one comes from truncating the distribution at \bar{k} . The column “p(top)” in Table 3 shows that this error can be neglected: only about 10^{-13} of all households are in the range between $0.9\bar{k}$ and \bar{k} . The second and more serious error arises insofar as the assumption in Equ. (27), which says that the density is constant within intervals, is not satisfied. To get an estimate of the order of magnitude of this error, I compute the steady state distribution both for $n_d = 1000$ and for $n_d = 5000$ and measure the difference. The table gives the maximum absolute distance between the two distribution functions, and the relative difference between the mean, standard deviation, skewness and kurtosis. We see that the mean is pinned down quite precisely, but higher moments become more and more difficult to pin down. Note that the difference in the invariant distributions is a tough criterion, since even minimal differences in the transition matrix can add up to sizable differences in the distribution over an infinite horizon. From this perspective, accuracy seems quite good.

5.3 The Approximation Error From State Space Reduction

In this section I take the Discrete Model **DM** with $n_d = 1000$ as the true model and investigate what is the error from using some approximate model **AM** with a reduced number of states. All exercises follow the same procedure:

1. Simulate the **DM** for $t = 1, \dots, T$, starting from the steady state distribution \mathbf{p}^* .
2. At each point t of the time series that results from this simulation, take the distribution \mathbf{p}_t as the starting point and simulate the **AM** for 10 periods, using the same shocks as the original simulation did. Compute the 1-step ahead and 10-step ahead forecast error, i.e., the difference in the value of \bar{K}_{t+1} and \bar{K}_{t+10} between the simulation of **DM** and the simulation of **AM**.

The focus on predicting the aggregate capital stock \bar{K}_t seems appropriate, because the only thing that matters for households are factor prices, which are determined by \bar{K}_t .

The above procedure implies that we compare the accuracy of the **AM** only at points in the state space that are “relevant” in the sense that they are reached in a long simulation. We cannot rule out that the approximation error is considerably higher at unlikely points in the state space.

5.3.1 Forecasting Future Aggregate Capital By Moments of the Distribution

For a model that is similar to ours, Krusell and Smith (1998) find that future values of the aggregate capital stock, and therefore of factor prices, can be very well predicted using only the following information: today’s aggregate level of capital, and the future realizations of the exogenous driving shocks. Table 4 investigates whether this is true in our model. We try to predict future values of capital by a VARX(1) model, where the endogenous variables are up to 4 moments of capital, and the exogenous variables are Z_t , Z_{t-1} and τ_t (not τ_{t-1} , since τ is i.i.d).

Table 4 presents standard deviations of forecast errors of horizon 1 year (equals 1 model period) and 10 years. The numbers reported are an average over 100 simulations of 1000 periods each.⁷ The parameters of the VARX(1) were estimated in sample. Out-of-sample prediction errors may be slightly higher. The first column of Table 4 indicates the driving shock: “Z” means that the technology shock was the only driving force, “tax” means that the tax shock was the only driving force, “Z,tax” means that both shocks are present. The second column gives the standard deviation of the one-step change in aggregate capital, $\sigma_K(1) \equiv stdev(K_t - K_{t-1})$. Columns 3–6 provide the standard errors of the VARX forecasts, from the case including only one moment up to the case of four moments. The standard errors are in relative terms, that means, divided by $\sigma_K(1)$. Columns 7–11 provide the analogous information for the 10-step-ahead forecasts.

If only technology shocks are present, the results confirm what earlier papers have reported: a forecast based only on the first moment (average capital) plus the exogenous

⁷For reasons of space, I do not report standard errors on these numbers. They are about 2% of the absolute value of the errors, such that the order of magnitude of the error is always well identified.

shock is already very precise. The one-step ahead prediction error is between a quarter of a percent and two percent of the period-by-period change. Ten years ahead, the forecast error is about the same fraction of the 10-period change. If we express this in terms of R^2 (forecast error variance relative to variance of aggregate capital), as is often done in the literature, it looks quite impressive. In the log-normal case with $\gamma = 1$, for example, the one-step R^2 is 0.9999998.

Things change once we allow for tax shocks that redistribute capital. If only tax shocks are operative, the variability of aggregate capital is much lower than with technology shocks. This is true in particular with the lognormal distribution, much less so with the Pareto distribution: with higher inequality, the distribution matters more for the dynamics of the aggregate, as we have explained in Section 5.1. In this case, a forecast based on one moment explains very little of the variation in aggregate capital. Even if we include 4 moments into the VAR, the forecast error is not always small in relative terms. If both shocks are operative, the variance of capital is big, and the accuracy is reduced by one or two orders of magnitude compared to the case with technology shocks only.

Notice that one should not interpret the numbers in Table 4 as the accuracy that a solution with the method of Krusell/Smith would achieve. It can better be seen as an upper bound: if the time series forecast of the simulation of the true model is not precise using n moments, a solution of the model where households base their forecasts on n moments only cannot be expected to be precise. In that case, a solution in a higher-dimensional state space is needed to get accurate results. I do not claim that the specification with the big tax shocks is realistic; the purposes of the exercise is only to create a laboratory in which one can test a method that uses a high-dimensional representation of the distribution.

5.3.2 The Accuracy of Spline Approximations of the Distribution

Figure 2 displays the 10-year forecast error from the Reduced Model (cf. Section 3.4) using spline bases. The solid line depicts the absolute forecast error, the dashed line the relative error, again defined as the absolute error divided by the standard deviation of $\bar{K}_{t+10} - \bar{K}_t$. Again, the numbers are the mean over 100 simulations of 1000 periods each, and Footnote 7 applies. The two isolated asterisks display the forecast error of the 1-moment and the 4-moment VARX, which are taken from Table 4. We see that the absolute forecast error for spline bases of dimension 100 is always around 10^{-6} or lower. Much more cannot be expected, since 10^{-8} is the square root of the machine precision, and the model solution contains a long series of complicated calculations. Moreover, the spline solutions are better than the 4-moment forecast by at least one order of magnitude, often more.

Finally, the first panel in Figure 2 contains a short circled line which displays the results of the state space reduction using principal components analysis. Using 5 components gives somewhat better results than 10; using only 3, accuracy would drop by an order of magnitude, which is not shown in the graph. However, using 10 or 30 principal components makes essentially no difference. The approach is successful, in the sense that it does better than spline bases if we are forced to work with a very reduced set of state variables. I do not want to stress this issue too much. It will not always work, and in some cases it will be

possible to find much better state aggregation approaches than the simple one proposed here. All this is better explored in the context of more interesting economic examples. I only wanted to make the point that the linear solution is a valuable tool to study the possibility of a state space reduction.

6 Conclusions

This paper has presented a new technique to solve models with a continuum of heterogeneous agents. The method allows to keep track of the dynamics of the cross-sectional distribution in a precise way. This was achieved by linearizing the model solution in the aggregate shocks, while maintaining the nonlinearity of the solution in the individual shocks. The next steps for future research are the following. First, to find out whether this is a suitable approach for studying asset pricing. Second, to compute a solution that is nonlinear in aggregate variables, in those cases where it is found that a low-dimensional representation of the distribution is sufficient. This is pursued in a follow-up paper that builds on the solution method in Reiter (2002).

References

- Algan, Y., O. Allais, and W. J. D. Haan (2006). Solving heterogeneous-agent models with parameterized cross-sectional distributions. manuscript.
- Bohacek, R. and M. Kejak (2005). Projection methods for economies with heterogeneous agents. CERGE-EI Working Paper 258.
- Carroll, C. D. (2001). A theory of the consumption function, with and without liquidity constraints. *Journal of Economic Perspectives* 15(3), 23–45.
- Carroll, C. D. (2004). Theoretical foundations of buffer stock saving. NBER Working Paper No. 10867.
- den Haan, W. J. (1997). Solving dynamic models with aggregate shocks and heterogeneous agents. *Macroeconomic Dynamics* 1, 355–86.
- Jin, H.-H. and K. L. Judd (2002). Perturbation methods for general dynamic stochastic models. Stanford University.
- Judd, K. L. (1998). *Numerical Methods in Economics*. Cambridge and London: MIT Press.
- Judd, K. L. (2004). Existence, uniqueness, and computational theory for time consistent equilibria: A hyperbolic discounting example. manuscript.
- Kim, J., S. Kim, and R. Kollmann (2005). Applying perturbation methods to incomplete market models with exogenous borrowing constraints. manuscript.
- Krusell, P. and A. Smith (2006). Quantitative macroeconomic models with heterogeneous agents. In R. Blundell, W. K. Newey, and T. Persson (Eds.), *Advances in economics*

- and econometrics: theory and applications, Ninth World Congress.* New York, NY: Cambridge University Press.
- Krusell, P. and J. Smith, Anthony A (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy* 106(5), 867–96.
- Miao, J. (2006). Competitive equilibria of economies with a continuum of consumers and aggregate shocks. *Journal of Economic Theory* 128, 274–298.
- Miranda, M. J. and P. L. Fackler (2002). *Applied Computational Economic and Finance*. MIT Press.
- Press, W., B. Flannery, S. Teukolsky, and W. Vetterling (1986). *Numerical Recipes*. Cambridge University Press.
- Preston, B. and M. Roca (2006). Incomplete markets, heterogeneity and macroeconomic dynamics. manuscript.
- Reiter, M. (2002). Recursive solution of heterogenous agent models. manuscript, Universitat Pompeu Fabra.
- Saez, E. (2001). Using elasticities to derive optimal income tax rates. *Review of Economic Studies* 68, 205–29.
- Sims, C. A. (2001). Solving linear rational expectations models. *Computational Economics* 20(1-2), 1–20.
- Stokey, N. L. and R. E. Lucas (1989). *Recursive Methods in Economic Dynamics*. Cambridge, MA: Harvard University Press.
- Virnik, E. (2006). An algebraic multigrid preconditioner for a class of singular M-matrices. manuscript, TU Berlin.

Table 1: Statistics of the income and stationary wealth distribution

$F(\xi)$	γ	Std ξ	Skew ξ	Kurt ξ	Mean k	Std k	Skew k	Kurt k	$p(\text{bottom})$
Logn	1	0.202	0.61	3.68	1.004	0.759	1.41	5.63	2.96
Pareto	1	1.373	27.90	1675.32	1.111	1.926	6.36	72.56	17.00
Logn	5	0.202	0.61	3.68	1.025	0.704	1.26	5.06	1.66
Pareto	5	1.373	27.90	1675.32	1.534	2.242	6.22	71.34	6.18

ξ : individual labor productivity

k : individual capital holdings

Std: standard deviation; Skew: skewness; Kurt: kurtosis

$p(\text{bottom})$: probability mass in the range $(0, 0.1)$.

Table 2: Times series variability of capital

$F(\xi)$	γ	Shocks Z		Shocks τ		Shocks Z, τ	
		Mean k	Var k	Mean k	Var k	Mean k	Var k
Logn	1	0.99	1.93	0.13	30.52	1.00	30.15
Pareto	1	0.90	0.90	0.86	10.44	1.23	10.47
Logn	5	1.65	4.23	0.41	27.97	1.68	27.38
Pareto	5	1.15	1.01	1.27	12.02	1.68	12.04

Mean k : 100 times standard deviation of log aggregate capital

Var k : 100 times standard deviation of log of cross-sectional variance of k

For the results of this table, $\rho_Z = \rho_\tau = 0$

Table 3: Approximation error from discretization

$F(\xi)$	γ	Abs. Euler resid.		Approx. error distribution k					
		max	ave	$p(\text{top})$	$\max_k \Psi_t(k)$	Mean k	Std k	Skew k	Kurt k
Logn	1	7.572e-5	1.296e-7	6.051e-14	4.650e-4	1.366e-5	2.049e-3	7.366e-3	1.034e-2
Pareto	1	6.974e-4	4.411e-7	1.016e-13	5.541e-5	5.490e-5	3.022e-4	4.063e-5	2.040e-3
Logn	5	9.790e-5	1.640e-7	2.376e-14	3.975e-4	5.785e-5	1.824e-3	6.436e-3	8.033e-3
Pareto	5	5.979e-4	3.358e-7	2.220e-15	3.447e-5	1.571e-4	4.140e-4	1.396e-4	1.710e-3

For the exact definition of the Euler residual, cf. the text

$p(\text{top})$ is the probability mass in the range $(0.9\bar{k}, \bar{k})$.

other statistics: difference between distribution with 1000 and with 5000 intervals

$\max_k \Psi_t(k)$: maximum difference in the distribution

Mean(k)–Kurt(k): relative difference in mean, standard deviation, skewness and kurtosis.

Table 4: Forecast errors of aggregate capital, from VARX(1)

	$\sigma_K(1)$	1 year ahead				$\sigma_K(10)$	10 years ahead			
		moments used					moments used			
		1	2	3	4		1	2	3	4
Lognormal income distribution, $\gamma = 1$										
Z	5.84e-3	2.30e-3	5.56e-4	3.77e-4	3.48e-4	3.26e-2	2.20e-3	5.00e-4	3.01e-4	2.81e-4
tax	1.40e-4	6.18e-1	6.10e-1	4.38e-1	1.48e-1	8.19e-4	6.29e-1	5.89e-1	2.65e-1	6.29e-2
Z,tax	5.84e-3	3.26e-2	1.89e-2	1.14e-2	8.03e-3	3.26e-2	2.83e-2	1.52e-2	7.92e-3	4.93e-3
Pareto income distribution, $\gamma = 1$										
Z	5.96e-3	2.81e-3	2.63e-3	2.27e-3	1.42e-3	3.33e-2	2.35e-3	1.95e-3	1.68e-3	8.76e-4
tax	1.85e-3	4.56e-1	1.58e-1	9.20e-2	6.21e-3	8.88e-3	4.79e-1	6.88e-2	4.90e-2	1.75e-2
Z,tax	6.25e-3	2.28e-1	8.55e-2	8.16e-2	8.07e-2	3.44e-2	1.87e-1	5.01e-2	4.74e-2	4.52e-2
Lognormal income distribution, $\gamma = 5$										
Z	6.67e-3	9.49e-3	2.11e-3	1.58e-3	1.20e-3	4.35e-2	1.09e-2	2.36e-3	1.74e-3	1.36e-3
tax	2.04e-4	8.18e-1	6.15e-1	1.12e-1	5.12e-2	1.57e-3	8.32e-1	4.98e-1	7.27e-2	2.66e-2
Z,tax	6.67e-3	4.30e-2	2.05e-2	1.20e-2	1.06e-2	4.35e-2	4.66e-2	2.04e-2	1.07e-2	9.03e-3
Pareto income distribution, $\gamma = 5$										
Z	7.95e-3	1.68e-2	9.16e-4	6.83e-4	4.10e-4	5.10e-2	1.80e-2	8.75e-4	6.84e-4	4.64e-4
tax	1.79e-3	7.29e-1	8.56e-2	4.47e-2	4.09e-2	1.23e-2	7.50e-1	4.80e-2	1.92e-2	1.46e-2
Z,tax	8.15e-3	2.22e-1	6.15e-2	5.59e-2	5.03e-2	5.23e-2	2.19e-1	4.66e-2	4.03e-2	3.83e-2

$\sigma_K(1)$: standard deviation of $\bar{K}_{t+1} - \bar{K}_t$

$\sigma_K(10)$: standard deviation of $\bar{K}_{t+10} - \bar{K}_t$

Forecast errors are relative to $\sigma_K(1)$ or $\sigma_K(10)$, respectively

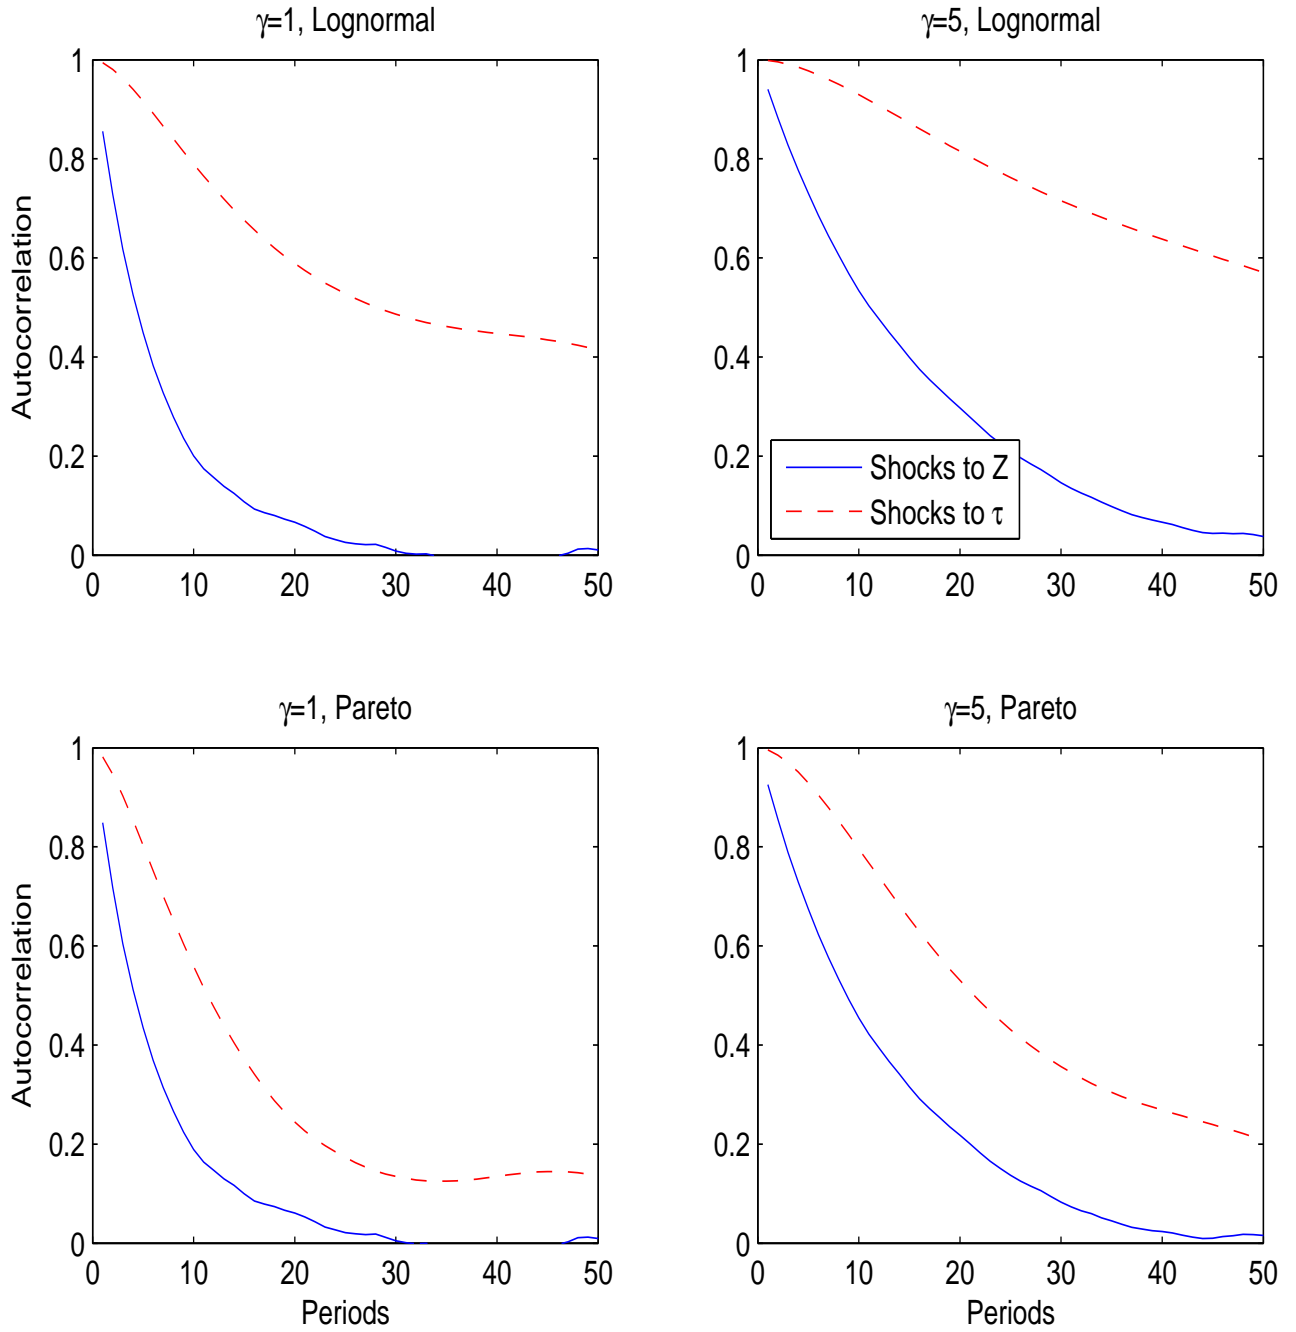


Figure 1: Autocorrelation function of aggregate capital, $\rho_Z = \rho_\tau = 0$

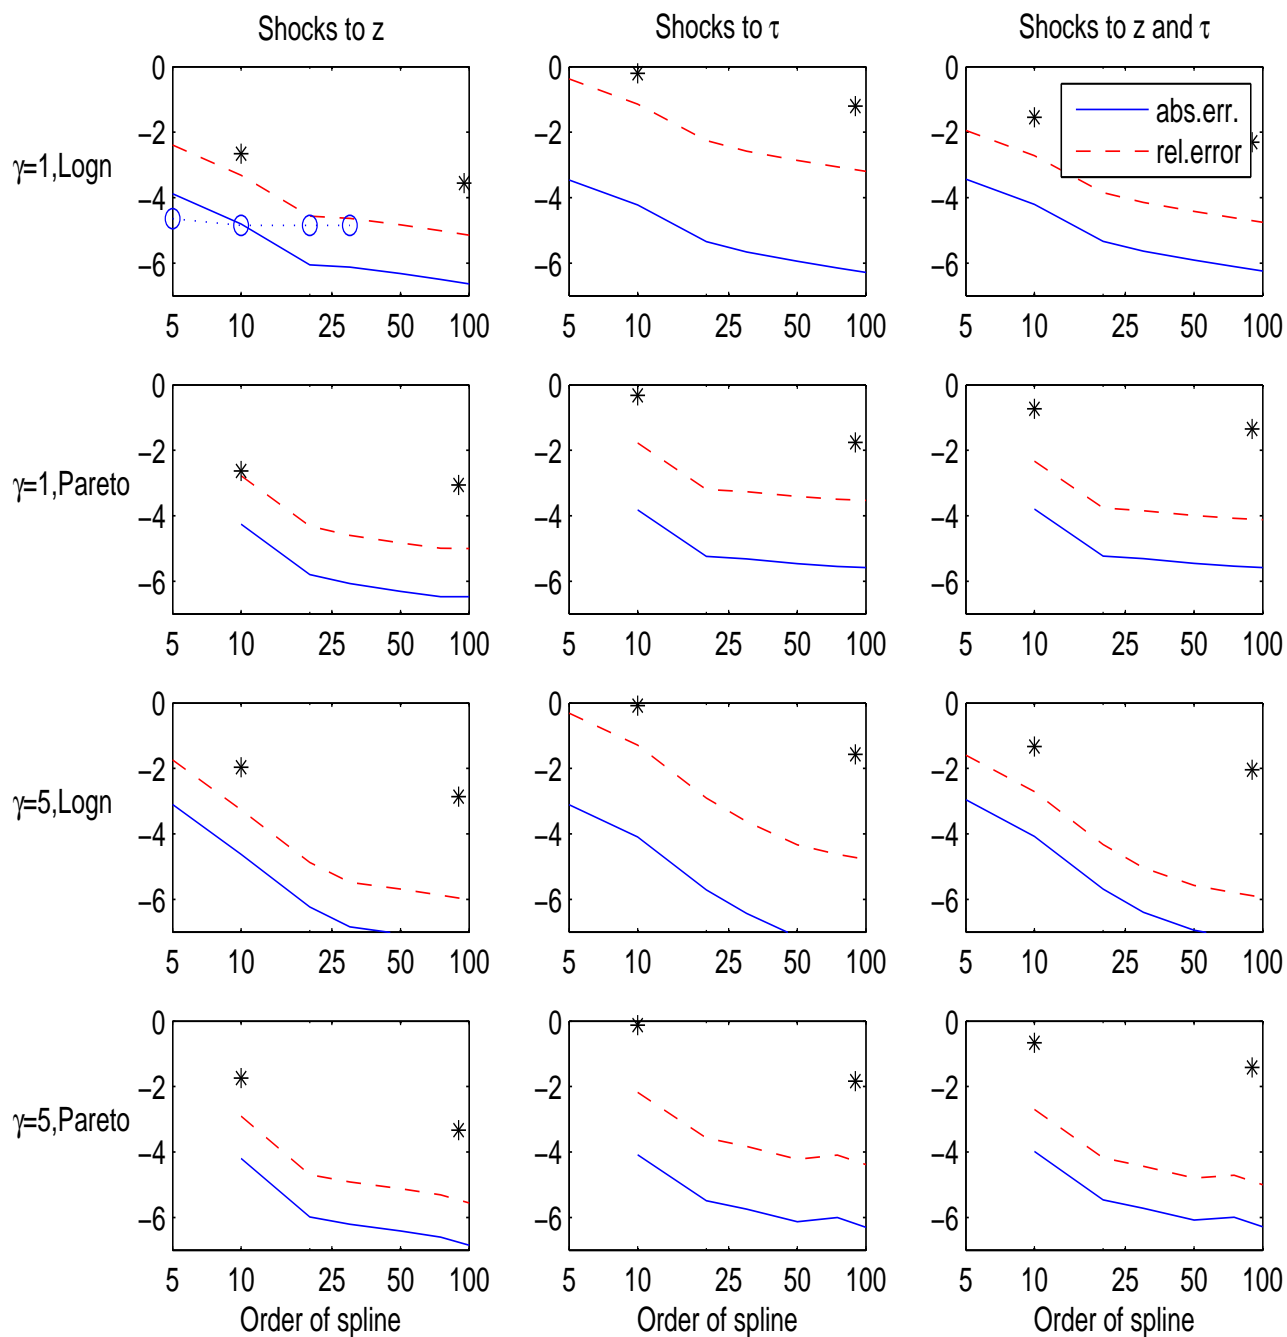


Figure 2: Log10 of 10-years forecast errors of aggregate capital, with spline bases