

Chapter 1: Heterogeneous Agent Models

Valerio Pieroni

Topics in Macroeconomics

University of Essex, PhD

I. Nonlinear systems

Let $f : \mathbb{R}^k \rightarrow \mathbb{R}^k$, suppose that we want to solve the nonlinear system $f(x) = 0$. We begin with the bisection algorithm. If $k = 1$ and f is continuous over the interval $[a, b]$, and $f(a), f(b)$ alternate the sign. Then, by the intermediate value theorem f must have at least one zero in $x \in [a, b]$. The bisection method start from an interval $[\underline{x}, \bar{x}]$ with $f(\underline{x}) < 0, f(\bar{x}) > 0$ and then evaluate f at the midpoint $x_n = (\underline{x} + \bar{x})/2$, if $f(x_n) = 0$ we are done, if $f(x_n) < 0$ set $\underline{x} = x_n$ and compute again the midpoint of this interval x_{n+1} , if $f(x_n) > 0$ set $\bar{x} = x_n$ and compute the midpoint x_{n+1} . We can stop when $|x_n - x_{n+1}| < \epsilon$ for $\epsilon \approx 0$ or $f(x_{n+1}) \leq \delta$ for $\delta \approx 0$. This procedure constructs successive smaller intervals containing a zero of f .

The Newton-Raphson method finds the roots of a nonlinear function using a sequence of linear approximations. We start by linearly approximating the function in a point x_n . The root of this line yields a new guess x_{n+1} . Convergence to the solution is guaranteed as long as the function is globally convex or globally concave. Suppose that $k = 1$, $f(x)$ is at least once differentiable. Solving $f(x_n) + (x_{n+1} - x_n)f'(x_n) = 0$ for x_{n+1} yields

$$x_{n+1} = x_n - \frac{f(x_n)}{f'(x_n)}.$$

From an initial guess x_n we can iterate until convergence, i.e. $|x_{n+1} - x_n| < \epsilon$, for ϵ small. In general, $f : \mathbb{R}^k \rightarrow \mathbb{R}^k$ the iteration scheme in matrix notation becomes

$$x_{n+1} = x_n - J_f(x_n)^{-1}f(x_n).$$

This requires to invert the Jacobian $J_f(x_n)$. When the Jacobian is computationally too demanding or ill-conditioned, there are alternative methods, called Quasi-Newton methods, that simplify its computation and provide approximations to it.

Given a transformation $g : \mathbb{R}^k \rightarrow \mathbb{R}^k$ we can rewrite root finding problems $f(x) = 0$ as fixed-point problems $x = g(x)$. This suggests the iteration $x_{n+1} = \theta x_n + (1 - \theta)g(x_n)$ until $|x_{n+1} - x_n| < \epsilon$. The parameter $\theta \in [0, 1]$ determines the size of each step affecting the speed and convergence of the algorithm.

The bisection method always converges but it is often slow. The Quasi-Newton algorithms tend to be faster but the convergence is not guaranteed. If a function is flat around the zero, e.g. x^6 , loose stopping rules can lead to convergence far from the true zero. Flat regions might also cause ill-conditioned matrices. Moreover, with irregular functions the

solution can be very sensitive to starting values due to the presence of multiple roots, and there are cases in which these methods produce cycles. Other pitfalls are related to rounding errors and scale problems. Fixed point schemes also do not guarantee convergence with the exception of existence results for the fixed points of contraction mappings.

II. Heterogeneous Agents and Distributions

Consider an intertemporal consumption problem where a_t denotes asset. Let w_t be the real wage and n_t labor supply. Households supply labor inelastically, i.e. $n_t = 1$ and labor income $w_t e_t$ is subject to an idiosyncratic risk e_t . The state space is $X = (A, E)$. We assume that e_t follows a two-state discrete Markov process. However, the code can be easily extended to allow for a $J > 2$ state Markov process. Since we assume incomplete markets the asset is not state-contingent. Given prices and states households solve

$$\begin{aligned} v(a, e) &= \max_{c, a'} u(c) + \beta \int_E v(a', e') dP_{e'|e}(e'|e), \\ \text{s.t. } c + a' &= (1 + r_t)a + w_t e, \\ a' &\geq -\phi \end{aligned} \tag{1}$$

Formally, $a_{t+1} = g_t^a(a_t, e_t)$ is the policy function of assets and $e_t : (\Omega, \mathcal{F}, P) \rightarrow (E, \mathcal{E})$ where $E = \{e_l, e_h\} \subset \mathbb{R}_+$ with conditional probability distribution $P_{e'|e}(\cdot | e = e_i) : \mathcal{E} \rightarrow [0, 1]$. Thus, the transition function $Q_t : (X, \mathcal{X}) \rightarrow [0, 1]$ of $\{a_t, e_t\}$ evaluated at (a', e') is

$$Q_t((a, e), (a', e')) = 1_{[g_t^a(a, e) = a']} P_{e'|e}(e'|e). \tag{2}$$

Intuitively, Q gives the probability of future states for someone with current states (a, e) . The function Q induces a sequence of distributions or probability measures D_t :

$$D_{t+1}(a', e') = \int_X Q_t((a, e), (a', e')) dD_t(a, e). \tag{3}$$

Let M be the set of probability measures on (X, \mathcal{X}) . It is convenient to use an implicit form for the previous law of motion $H_t : M \rightarrow M$,

$$D_{t+1} = H_t(D_t).$$

There is a continuum $\mathcal{H} = [0, 1]$ of ex-ante identical agents subject to idiosyncratic labor income shocks. Hence, the joint probability distribution $D_t(a_t, e_t)$ gives us the cross-sectional distribution of assets and income risk in period t . To begin with, we consider the steady state of the model taking the prices w and r as given. In Section III. we consider the case in which prices are determined endogenously in general equilibrium and show how to compute the transition dynamics of the model.

A. The Endogenous Grid Method

To solve the household optimization problem we can use Value Function Iteration (VFI). However, the endogenous grid method provides a more efficient alternative. As in the case of discrete VFI, we discretize the state space A using a grid G_A of points in A . However, rather than using a grid $G_A \in A$ over a we use a grid $G_A \in A$ over a' . This allows us to solve the Euler equation analytically. As a result, this substantially speeds up the computations because we can avoid the inner loop that solves the Euler equation or the maximization in the Bellman equation numerically. We set $w = 1$ so that $y = e$. The algorithm is the following:

1. Fix a grid over assets $G_A = \{a_1, \dots, a_I\}$ and income shocks $G_Y = Y = \{y_1, \dots, y_J\}$.
2. Guess a decision rule $c_n(a_i, y_j)$, e.g. $a'_n(a_i, y_j) = 0, c_n(a_i, y_j) = (1 + r)a_i + y_j$.
3. For any pair $(a_i, y_j) \in G_A \times G_Y$ compute the right-hand-side of the Euler equation

$$RHS_{ij} = \beta(1 + r) \sum_{j'=1}^J P(y_{j'}|y_j) u'(c_n(a'_i, y_{j'})).$$

4. Compute $c(a'_i, y_j) = (u')^{-1}(RHS_{ij})$ and $a(a'_i, y_j) = (1 + r)^{-1}(c(a'_i, y_j) + a'_i - y_j)$.
5. For each j invert the mapping a to go from $(a'_i, a(a'_i))$ to $(a(a'_i), a'_i)$.
6. For each j interpolate $(a(a'_i), a'_i)$ over $a_i \in G_A$ to get $a'_{n+1}(a_i, y_j)$ and $(a(a'_i), c(a'_i, y_j))$ over $a_i \in G_A$ to get $c_{n+1}(a_i, y_j)$.
7. If $a_i < a(a'_1, y_j)$ the agent will be constrained in the next period and we cannot use the Euler equation. Hence, set $a'_{n+1} = a_1$ and $c_{n+1}(a_i, y_j) = (1 + r)a_i + y_j - a_1$.
8. Iterate from step 3 until convergence $\max_{ij} |c_{n+1} - c_n| < \varepsilon$.

Note that $a(a'_i, y_j)$ are the assets today that will lead the consumer to have a'_i assets tomorrow given the shock today y_j . This function is not necessarily on the grid G_A and is an endogenous asset grid that changes in each iteration. In step 6 we want to move the optimal decisions from the endogenous grid to the exogenous grid G_A , if needed we can always extrapolate. The algorithm can be extended to include a labor supply decision and finite horizon to solve life-cycle models. Note, that one can also apply this method with minimum modifications to solve the first-order conditions of the recursive formulation of the household optimization problem instead of the Euler equation of the sequential household problem as we did.

B. Stationary distribution

Once we have the policy functions we can check that g^a intersects the 45-degree line for a large enough to rule out explosive wealth dynamics. We can always adjust the upper bound of the grid a_I to make sure that we are covering the relevant section of the state space A . Then, we can compute the stationary distribution. First, we compute the transition probabilities given by Q from pairs (a, e) to pairs (a', e') . Since we discretized the state space we are working on a grid $G_A \times G_E$ and we can collect these probabilities in a $IJ \times IJ$ matrix A . Note here that the asset policy function $a' = g^a(a, e)$ will take us off the grid in the next period. We can use a linear interpolation. In particular, we assume that households go to the two nearest grid points so that on average households choose the right asset value. To do so we solve $pa_i + (1 - p)a_{i+1} = a'$ for p . Second, we solve the discrete Chapman-Kolmogorov equation, i.e. $D_{t+1} = A'D_t$ where D_t is a $IJ \times 1$ vector. There are three equivalent ways to do so. Since in the steady state $D = H(D)$ we can use a fixed point iteration scheme

$$D^{n+1} = A'D^n,$$

and iterate until convergence $|D^{n+1} - D^n| < \varepsilon$. We can use the eigenvalue problem

$$A'D = \lambda D,$$

the eigenvector associated to the eigenvalue $\lambda = 1$ is the stationary distribution D . We can solve this easily as it is very likely that the programming language you are using already has routines for it. We should rescale D to make sure that it adds up to 1. Alternatively, simulate a large number of households say 10,000 initialize each individual at (a_0, e_0) and use g^a and a random number generator to replicate the Markov process $\{e_t\}$ and generate the joint Markov process $\{a_t, e_t\}$. In each period compute a set of cross-sectional moments m_t for the distribution of assets like mean, variance, and quantiles. When $m_t \approx m_{t+1}$ we can stop. In this case, the distribution has converged.

Once we have the policy functions and the stationary distribution we have solved the heterogeneous agent model and we can compute aggregate consumption and wealth

$$A = \sum_{(a,e)} g^a(a, e) D(a, e),$$

$$C = \sum_{(a,e)} g^c(a, e) D(a, e),$$

As well as marginal distributions $D_a(a_t)$ and $D_e(e_t)$, Gini coefficients, correlations, and expectation functions such as $E_h^x(a, e) = \mathbb{E}[x(a_{t+h}, e_{t+h}) | a_0 = a, e_0 = e]$.

III. The Workhorse Heterogeneous Agent (HA) model

A. Description of the model

In this section we study the canonical heterogeneous-agent model with capital and no aggregate risk (Aiyagari (1994), Imrohoroglu (1992), Huggett (1993), Krusell and Smith (1998)). Given prices and states households solve the same consumption problem as before

$$v(a, e) = \max_{c, a'} u(c) + \beta \int_E v(a', e') dP_{e'|e}(e'|e), \quad (4)$$

$$\text{s.t. } c + a' = (1 + r_t)a + w_t e, \quad (5)$$

$$a' \geq -\phi$$

The law of motion of the distribution of idiosyncratic states D_t is given by

$$D_{t+1}(a', e') = \int_X Q_t((a, e), (a', e')) dD_t(a, e). \quad (6)$$

Firms operate in a competitive sector and produce the final consumption good using labor inputs L_t and capital K_t . Marginal pricing implied by profit maximization yields

$$Y_t = K_t^\alpha L_t^{1-\alpha}, \quad (7)$$

$$r_t + \delta = \alpha K_t^{\alpha-1} L_t^{1-\alpha}, \quad (8)$$

$$w_t = (1 - \alpha) L_t^{-\alpha} K_t^\alpha. \quad (9)$$

In each period labor and asset markets clear:

$$L_t = \int_X e_t dD_t(a_t, e_t) = \sum_{e_t \in E} \int_0^\infty e_t f_t(a_t, e_t) da_t, \quad (10)$$

$$K_t = \int_X a_t dD_t(a_t, e_t) = \sum_{e_t \in E} \int_0^\infty a_t f_t(a_t, e_t) da_t, \quad (11)$$

where $f_t(a_t, e_t)$ is the probability function of a_t, e_t associated to D_t . This is a mixed probability function as one random variable is discrete e_t and the other a_t is continuous. Since labor is supplied inelastically labor supply is exogenous. The resource constraint is

$$C_t + K_{t+1} = Y_t + (1 - \delta)K_t,$$

where aggregate consumption is given by

$$C_t = \int_X c_t(a_t, e_t) dD(a_t, e_t).$$

B. Stationary equilibrium

In equilibrium households decide $\{c_t, a_t\}$ solving (4), (5) given prices $\{w_t, r_t\}$ and initial conditions a_0, e_0 , firms demand $\{K_t, L_t\}$ and produce $\{Y_t\}$ according to (7), (8), (9) given prices $\{w_t, r_t\}$, $\{D_t\}$ follows the law of motion (6), and prices $\{w_t, r_t\}$ are such that markets (10), (11) clear. This is a nonlinear dynamic system with 8 equations and 8 endogenous variables.

In the presence of idiosyncratic risk, we need $\beta(1+r) < 1$ else there is no stationary equilibrium as the sequence of consumption and assets are unbounded. This result has an intuitive economic interpretation: the presence of idiosyncratic income risk implies a precautionary saving motive leading to more capital accumulation and a lower equilibrium rate than under the complete markets benchmark $\beta(1+r_{cm}) = 1$.

Example. Figure 1 shows the consumption and saving policy functions of the canonical HA model where e_t follows a two-state Markov process. The consumption function is concave around the borrowing limit ϕ , in this example $\phi = 0$. One important statistic in this class of models is the Marginal Propensity to Consume (MPC), that is the fraction of a windfall income or wealth gain that is consumed within a given period. This statistic summarizes how responsive is households' expenditure to temporary income changes. The concavity of the consumption function implies a high MPC for constrained agents and low-wealth unconstrained agents close to the borrowing limit (see Chapter 3 for more details on this point). On the other hand, wealthy households are well insured against income shocks and the consumption function is linear in this region of the state space.




Figure 1: Consumption and saving policy functions.

The saving policy function $s(a, e) := a'(a, e) - a$ shows that households with a low income realization reduce their wealth $s(a, e) < 0$ while households with a high income

realization accumulate wealth $s(a, e) > 0$ until they reach their saving target $a^* : s(a^*, e) = 0$. This saving function implies that the support of the wealth distribution is bounded. Alternatively, one could represent saving as a stock a' and use a phase diagram (a', a) to make the same point.

C. Solving the model using the sequence space Jacobian

We solve the model in sequence space as in [Auclert, Bardóczy, Rognlie, and Straub \(2021\)](#). In particular, we need to solve a dynamic system of nonlinear equations in sequence space given by

$$M(X, Z) = 0,$$

where X is a sequence of endogenous variables and Z is the time path of the exogenous variables. Often we can reduce this problem to a few unknown sequences U such that

$$H(U, Z) = 0,$$

in this case, once we obtain the unknowns we can recover other structural variables using the equations of the model, namely $X = F(U, Z)$. To solve the model globally we can use a Quasi-Newton algorithm. Take a linear approximation of the model $H(U, Z) = 0$ at an initial point (U^n, Z) this yields $H(U^n, Z) + H_U(U^n, Z)(U - U^n) = 0$. Replacing $H_U(U^n, Z)$ with $H_U(U_{se}, Z_{se})$ and solving for $U^{n+1} := U$ yields

$$U^{n+1} = U^n - H_U(U_{se}, Z_{se})^{-1} H(U^n, Z).$$

Alternatively, we can take only one linear approximation of the model around the steady state (U_{se}, Z_{se}) for a local solution. At the steady state $H(U_{se}, Z_{se}) = 0$. Totally differentiating yields $H_U(U_{se}, Z_{se})dU + H_Z(U_{se}, Z_{se})dZ = 0$. Rearranging terms yields

$$dU = -H_U(U_{se}, Z_{se})^{-1} H_Z(U_{se}, Z_{se})dZ.$$

In both cases we need the Jacobian matrix $H_U(U_{se}, Z_{se})$. In discrete time with T periods and n unknowns this is a $nT \times nT$ matrix.

To compute the model Jacobian $H_U(U_{se}, Z_{se})$ we divide the model into blocks. Each block is a set of m input variables, o output variables and structural equations. We order the blocks using a Directed Acyclical Graph (DAG). The graph should have as many outputs or targets H as inputs U . Moreover, you cannot have cycles in the graph: U cannot be the output of any intermediate block and H cannot be the input of any intermediate block. Note that there might be many DAG representations of a given model. All DAGs of a model deliver the same solution, finding an efficient DAG becomes easier with practice. Then, we obtain H_U by forward accumulation of the partial Jacobians along the DAG.

Example. Consider the canonical HA model with a TFP shock $Y_t = Z_t K_t^\alpha L_t^{1-\alpha}$ where Z_t follows an AR(1) process and $Z = 1$ at the steady state. Let $C_t(\{w_t, r_t\})$, $A_t(\{w_t, r_t\})$ denote respectively aggregate consumption and aggregate household wealth from the heterogeneous agent block of the model. To solve the model we use the following DAG.




Figure 2: Computations of the standard Heterogeneous Agent model

Each node in Figure 2 is a block of the model. Note that labor supply is exogenous and given by $L_t = \sum_e e_t f_e(e_t)$ where f_e is the marginal probability function of the income risk process. So, we have one unknown sequence $\{K_t\}$ and one target given by the market clearing condition $K_t = A_t, \forall t = 0, 1, 2, \dots, T$ where T is a truncation horizon. T must be large enough that the economy is back at the steady state. We compute the Jacobian $H_K(K_{se}, Z_{se})$ using the partial Jacobian of each block and the chain rule to forward accumulate these matrices as follows. In the first block we use Equations (7), (8), (9) to obtain the block's output variables w, r given the block's inputs K, L, Z . Specifically, in this block we compute the $T \times T$ partial Jacobians $\mathcal{J}^{w,K}, \mathcal{J}^{w,Z}, \mathcal{J}^{r,K}, \mathcal{J}^{r,Z}$. Each column s of these matrices is the response of an output variable to a one-time shock of size $h = 0.0001$ to an input variable while keeping the other inputs at the steady state, e.g. $\mathcal{J}_{t,s}^{w,K} := (w_t - w)/h$ is the partial response of real wages when aggregate capital is increased above the steady state level by h at time s while keeping the other inputs at the steady state level. In the second block we compute again the partial jacobians $J^{A,r}, J^{A,w}$ as before. Then, we compute the total Jacobians J with respect to the initial inputs, i.e. the unknowns and the shocks, using the chain rule $J^{A,K} = \mathcal{J}^{A,r} J^{r,K} + \mathcal{J}^{A,w} J^{w,K}$ and $J^{A,Z} = \mathcal{J}^{A,r} J^{r,Z} + \mathcal{J}^{A,w} J^{w,Z}$. Note that partial derivatives \mathcal{J} and total derivatives J are the same in the first block of the model. Moreover, $\mathcal{J}^{K,K} = J^{K,K} = I_T$ where I_T is the $T \times T$ identity matrix. Finally, in the last block we proceed as before and compute $\mathcal{J}^{H,K}, \mathcal{J}^{H,A}$ and $H_K := J^{H,K} = \mathcal{J}^{H,A} J^{A,K} + \mathcal{J}^{H,K}$ and $H_Z := J^{H,Z} = \mathcal{J}^{H,A} J^{A,Z}$. These computations can be easily automatized. However, for the basic model it is instructive to build the Jacobian H_U "by hand".

Note that to compute the Jacobian H_U and solve the model. We need to solve the heterogeneous agent block of the model over time. To this end, we solve the household problem backward in time starting at $t = T$ from a guess of the policy functions in $c_T = c, a_T = a'$. Then, we solve the distribution forward in time starting at $t = 0$ from the steady state distribution $D_0 = D$. This implies that to compute each column of the Jacobians we need to iterate twice over the time dimension. Usually, $T = 300$ so this is a bottleneck for the numerical solution. One could speed up the algorithm using one single backward iteration for each Jacobian exploiting the fact that the policy functions are purely forward-looking.

D. Productivity shocks

Figures 3 plot the dynamics of the model after an unexpected 1% increase in productivity Z_t with autoregressive coefficient $\rho_z = 0.8$. This is a local solution around the steady state of the model. The TFP shock increases output and the factor prices stimulating aggregate consumption and firms' demand for labor and capital. Higher real interest rates provide an incentive for saving while higher real wages increase consumption of high-MPC households. Over time the productivity shocks fade away and capital depreciates. This generates a hump-shaped response of households' assets and aggregate consumption.




Figure 3: Impulse response functions to an increase in productivity

Note: The response of the variables are in percentage deviations from the steady state. The response of the real rate is in percentage points.

References

- Aiyagari, Rao (1994). “Uninsured Idiosyncratic Risk and Aggregate Savings”. In: *The Quarterly Journal of Economics* 109 (3), pp. 659–684.
- Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub (2021). “Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models”. In: *Econometrica* 89 (5), pp. 2375–2408.
- Huggett, Mark (1993). “The risk-free rate in heterogeneous-agent incomplete-insurance economies”. In: *Journal of Economic Dynamics and Control* 17 (5), pp. 953–969.
- Imrohoroglu, Ayse (1992). “The welfare cost of ináation under imperfect insurance”. In: *Journal of Economic Dynamics and Control* 16 (1), pp. 79–91.
- Krusell, Per and Anthony Smith (1998). “Income and Wealth Heterogeneity in the Macroeconomy”. In: *Journal of Political Economy* 106 (5), pp. 867–896.