Lecture III
Solving Heterogeneous Agent General Equilibrium Models with Idiosyncratic Uncertainty

Adam Hal Spencer

The University of Nottingham

Applied Computational Economics 2020

Roadmap

Motivation

- We're all used to seeing the representative agent framework.
- Such a setup need not assume that there is literally only one agent, but rather a degenerate distribution across agents.
- I.e. a representative agent setup abstracts from thinking about cross-sectional dispersion.

Motivation

- But the cross-section can have implications for aggregates!
- Especially in policy settings.
- E.g. HANK (heterogeneous agent new Keynesian) models.
- Kaplan, Moll and Violante (2018, AER), "Monetary Policy According to HANK"...

Motivation

We revisit the transmission mechanism from monetary policy to household consumption in a Heterogeneous Agent New Keynesian (HANK) model. The model yields empirically realistic distributions of wealth and marginal propensities to consume because of two features: uninsurable income shocks and multiple assets with different degrees of liquidity and different returns. In this environment, the indirect effects of an unexpected cut in interest rates, which operate through a general equilibrium increase in labor demand, far outweigh direct effects such as intertemporal substitution. This finding is in stark contrast to small- and medium-scale Representative Agent New Keynesian (RANK) economies, where the substitution channel drives virtually all of the transmission from interest rates to consumption. Failure of Ricardian equivalence implies that, in HANK models, the fiscal reaction to the monetary expansion is a key determinant of the overall size of the macroeconomic response.

Motivation

- Yours truly also wrote a job market paper (a long time ago) looking at the effect of policy changes in the face of firm heterogeneity.
- For the particular policy I studied, the change has close to no effect on the economy in a representative firm framework.
- Does in fact have a quantitatively significant impact with heterogeneous firms and selection effects.
- Changes in the cross-section can aggregate to affect the macroeconomy!

Motivation

- In this course, we'll focus mainly on studying heterogeneity on the household side.
- You'll think about firms in your other computational masterclass.

Roadmap

Motivation

- Why does heterogeneity matter?
- A classic issue it allows us to think about is inequality.
- What do we need in a model of household saving to get at inequality?
- Start with a baseline complete markets model then see what, (if anything), needs to be added to allow us to study this issue.

Motivation

- Assume a population of households indexed by $j \in \{1, 2, ..., J\}$.
- Two time periods $t \in \{0, 1\}$.
- Time $t = 0$ state known, but $t = 1$ state unknown. State space is Ω, known to all households as is the distribution across states.
- Probability of state $\omega \in \Omega$ denoted by $\pi(\omega)$.
- Household j receives endowment of $y_t^j(\omega)$ in $t \in \{0, 1\}$.
- Households consume each period and purchase a portfolio of state-contingent claims: price of a claim for state ω is denoted by $\varphi(\omega)$.
- All households have the same period utility function $u(c_t^j)$ and discount factor $\beta \in [0, 1]$.

Motivation

- Household's problem

$$\max_{\{c_0^j, \{b_1^j(\omega), c_1^j(\omega)\}_{\omega\in\Omega}\}} u(c_0^j) + \beta \sum_{\omega\in\Omega} \pi(\omega) u(c_1^j(\omega))$$

subject to

$$\sum_{\omega\in\Omega} b_1^j(\omega)\varphi(\omega) + c_0^j = y_0^j$$

$$c_1^j(\omega) = y_1^j(\omega) + b_1^j(\omega) \;\; \forall \omega \in \Omega$$

where notice that the second constraint holds for all states of the world at $t = 1$.

Motivation

- Lagrangian

$$\mathscr{L} = u\left(c_0^j\right) + \beta \sum_{\omega\in\Omega} \pi(\omega) u\left(c_1^j(\omega)\right)$$
$$+ \lambda_0^j \left[y_0^j - \sum_{\omega\in\Omega} b_1^j(\omega)\varphi(\omega) - c_0^j\right] + \sum_{\omega\in\Omega} \lambda_1^j(\omega)\left[y_1^j(\omega) + b_1^j(\omega) - c_1^j(\omega)\right]$$

Motivation

- FOCs

$$\frac{\partial \mathscr{L}}{\partial c_0^j} = u'(c_0^j) - \lambda_0^j = 0$$
$$\frac{\partial \mathscr{L}}{\partial b_1^j(\omega)} = -\lambda_0^j \varphi(\omega) + \lambda_1^j(\omega) = 0$$
$$\frac{\partial \mathscr{L}}{\partial c_1^j(\omega)} = \beta \pi(\omega) u'(c_1^j(\omega)) - \lambda_1^j(\omega) = 0$$

Motivation

- See then that the A-D securities are priced as

$$\varphi(\omega) = \beta\pi(\omega)\frac{u'(c_1^j(\omega))}{u'(c_0^j)}$$

Motivation

- Equilibrium requires that A-D securities are held in zero net supply.
- If $J = 1$, (i.e. only one agent in the economy), what does this mean?
- Autarky! There's no trading of anything.
- Prices $\varphi(\omega)$ will adjust such that the sole agent is incentivised to not hold any A-D securities.
- **Motivation 1:** this is really stupid; we need multiple non-identical agents to get trading in these securities markets.
- Is heterogeneity enough though...?

Motivation

- Now come back to the general case with $J > 1$.
- For any two agents, call them $j, k \in \{1, 2, ..., J\}$ with $j \neq k$, the Euler equations imply

$$\frac{u'(c_1^j(\omega))}{u'(c_0^j)} = \frac{u'(c_1^k(\omega))}{u'(c_0^k)}$$

which says we have full risk sharing.

- This holds across all states $\omega \in \Omega$.
- With complete markets, the marginal rates of substitution are equalised across households.
- With CRRA preferences, this amounts to saying that consumption growth must be equalised across households for all states.

Motivation

- With homogeneous CRRA preferences, each household ends up consuming a constant fraction of the aggregate endowment.
- No agent bears any idiosyncratic risk.
- In fact, with complete asset markets and heterogeneous agents, we can find an equivalent economy with a representative agent that generates the same outcomes.
- The key is to construct welfare weights appropriately for the different agents.
- Just like a single agent consumes the entire endowment.

Motivation

- If no agent bears his own idiosyncratic risk, how the hell can we say anything meaningful about inequality?
- **Motivation 2:** we also need incomplete markets to study inequality.

Motivation

- Assume now that instead of Arrow-Debreu securities, households can only save through a riskless bond that pays $R > 1$ (exogenous) gross interest at $t = 1$.
- Denote their holdings of this asset by a_1^j.

Motivation

- Household's problem

$$\max_{\{c_0^j, a_1^j, \{c_1^j(\omega)\}_{\omega \in \Omega}\}} u(c_0^j) + \beta \sum_{\omega \in \Omega} \pi(\omega) u(c_1^j(\omega))$$

subject to

$$a_1^j + c_0^j = y_0^j$$
$$c_1^j(\omega) = y_1^j(\omega) + R a_1^j \;\; \forall \omega \in \Omega$$

Motivation

- Lagrangian

$$\mathscr{L} = u\left(c_0^j\right) + \beta \sum_{\omega \in \Omega} \pi(\omega) u\left(c_1^j(\omega)\right) + \lambda_0^j \left[y_0^j - a_1^j - c_0^j\right] + \sum_{\omega \in \Omega} \lambda_1^j(\omega) \left[y_1^j(\omega) + a_1^j - c_1^j(\omega)\right]$$

Motivation

- FOCs

$$\frac{\partial \mathscr{L}}{\partial c_0^j} = u'(c_0^j) - \lambda_0^j = 0$$

$$\frac{\partial \mathscr{L}}{\partial a_1^j} = -\lambda_0^j + \sum_{\omega \in \Omega} R \lambda_1^j(\omega) = 0$$

$$\frac{\partial \mathscr{L}}{\partial c_1^j(\omega)} = \beta \pi(\omega) u'(c_1^j(\omega)) - \lambda_1^j(\omega) = 0$$

Motivation

- Euler equation now given by

$$
\begin{aligned}
u'(c_0^j) &= \beta R \sum_{\omega \in \Omega} \pi(\omega) u'(c_1^j(\omega)) \\
&= \beta R \mathbb{E}_0[u'(c_1^j)]
\end{aligned}
$$

- No longer have this state-by-state equivalence across households.
- The marginal rates of substitution are only equal in expectation.
- Now we're in a position to start talking about inequality.

Motivation

- These heterogeneous agent models are often referred to as incomplete markets models.
- Agents are assumed to be ex-ante identical but heterogeneous ex-post.
    - Sometimes called Bewley models.
    - Bewley (1977), "The Permanent Income Hypothesis: a Theoretical Formulaiton", *Journal of Economic Theory*, 16(2), pp. 252–292.

Motivation

- Bewley incomplete markets models made quantitative by Hugget and Aiyagari:
  - Huggett, (1993), "The Risk-Free Rate in Heterogeneous Agent Incomplete-Insurance Economies", *Journal of Economic Dynamics and Control*, 17, pp. 953–969.
  - Aiyagari (1994), "Uninsured Idiosyncratic Risk and Aggregate Saving", *Quarterly Journal of Economics*, 109(3), pp. 659–684.

Roadmap

Borrowing Constraints

- Before we jump-into one of these GE idiosyncratic agent models, let's think a bit more about constraints.
- Take the incomplete markets model from the last section: recall that households could save through $a_1^j \geqslant 0$.
- What if we also allow them to borrow?
- I.e. $a_1^j < 0$.
- Assume gross interest of $R > 1$ (still exogenous).

Borrowing Constraints

- There must be some limit, otherwise these self-obsessed agents will start running schemes with infinite borrowing!
- Standard is to assume some exogenous debt limit

$$a_1^j \geqslant \underline{a}$$

for some $\underline{a} \leqslant 0$.

Borrowing Constraints

- Now their problem becomes, (dropping the statewise notation now)

$$\max_{\{c_0^j, a_1^j, c_1^j\}} u(c_0^j) + \beta \mathbb{E}_0[u(c_1^j)]$$

subject to

$$a_1^j + c_0^j = y_0^j$$
$$c_1^j = y_1^j + Ra_1^j$$
$$a_1^j \geqslant \underline{a}$$

Borrowing Constraints

- Lagrangian

$$\begin{aligned}\mathscr{L} &= u(c_0^j) + \beta \mathbb{E}_0[u(c_1^j)] \\ &+ \lambda_0^j \left[y_0^j - a_1^j - c_0^j\right] + \lambda_1^j \left[y_1^j + Ra_1^j - c_1^j\right] + \mu[a_1^j - \underline{a}]\end{aligned}$$

where $\mu \geqslant 0$ is the Lagrange multiplier on the borrowing constraint.

Borrowing Constraints

- FOCs

$$\frac{\partial \mathscr{L}}{\partial c_0^j} = u'(c_0^j) - \lambda_0^j = 0$$
$$\frac{\partial \mathscr{L}}{\partial a_1^j} = -\lambda_0^j + R\lambda_1^j + \mu = 0$$
$$\frac{\partial \mathscr{L}}{\partial c_1^j} = \beta \mathbb{E}_0[u'(c_1^j)] - \lambda_1^j = 0$$

Borrowing Constraints

- Euler equation

$$u'(c_0^j) = R\beta\mathbb{E}_0[u'(c_1^j)] + \mu$$

or can alternatively think of this as an Euler inequality

$$u'(c_0^j) \geqslant R\beta\mathbb{E}_0[u'(c_1^j)].$$

- We call these constraints occasionally binding, since $\mu = 0$ sometimes, meaning we get our traditional Euler equation back.

Borrowing Constraints

- If we extend to an infinite horizon (still discrete time) setup, we can write

$$u'(c_t^j) \geqslant R\beta\mathbb{E}_t[u'(c_{t+1}^j)].$$

- Define $\theta_t = \beta^t R^t u'(c_t) \geqslant 0$. Then our Euler inequality says

$$\theta_t \geqslant \mathbb{E}_t[\theta_{t+1}]$$

Martingales

- A martingale is a stochsatic process $\{X_t, t \geqslant 1\}$ such that
  - $\mathbb{E}[|X_t|] < \infty$ and
  - $\mathbb{E}[X_{t+1}|X_1, X_2, ..., X_t] = X_t$.
- A submartingale is such that

$$\mathbb{E}[X_{t+1}|X_1, X_2, ..., X_t] \geqslant X_t$$

where here notice that $-\theta_t$ from the previous slide is a submartingale.

Martingales

- **Theorem (Martingale Convergence)**: let $\{X_t\}$ be a submartingale. If $K = \sup_t \mathbb{E}(|X_t|) < \infty$ then $X_t \to X$ with probability one where X is a random variable such that $\mathbb{E}(|X|) \leqslant K$.

Martingales

- Since $-\theta_t$ is a submartingale, follows that $\theta_t \to \bar{\theta}$.
- See that if $\beta R > 1$ then since $\lim_{t\to\infty} \beta^t R^t = \infty$, must be that $\lim_{t\to\infty} u'(c_t) = 0$ with probability 1 so $c_t \to \infty$. Implies that $a_t \to \infty$ if $\beta R > 1$.
- Similar arguments hold for $\beta R = 1$.
- Thus it must be that $\beta R < 1$ in the incomplete markets case.
- Agents build-up a stock of assets to insure against their income risk.
- Drives-down the equilibrium interest rate relative to the complete markets case.

Borrowing Constraints: Implementation

- How to we account for the constraint in our computations?
- These constraints look intimidating. They're no big deal though...
- ...provided that you're coding the model up yourself.
- These constraints impede one's ability to use perturbation methods.
- Although I think Dynare has a toolkit called *occbin*, which can account for them...
- ... this is incredibly "black-boxy" though.

Borrowing Constraints: Implementation

- When we code-up the model ourselves, accounting for an occasionally binding constraint just involves an extra 3–4 lines of code.
- Just manually adjust the policy function such that the constraint is accounted for.

Borrowing Constraints: Implementation

- In general, assume a household's control x' must be above some variable $\underline{x}(x)$.
- I.e. a function of your current state x.

Value Function Iteration and Constraints

- Again, augment the VFI algorithm. Say we use gridsearch.
- If our value function is of the form

$$V(x,a) = \max_{x' \geqslant \underline{x}(x)} u(x,x',a) + \beta \mathbb{E}_{a'}[V(x',a')]$$

- Say we have an initial guess of $V_0(x,a)$.
- Find all the values associated with each candidate $x' \in \mathscr{X}$ where

$$\begin{aligned}\tilde{V}_1(x,x',a) &= u(x,x',a) + \beta \mathbb{E}_{a'}[V_0(x',a')] \text{ if } x' \geqslant \underline{x}(x) \\ &= -\infty \text{ if } x' < \underline{x}(x)\end{aligned}$$

  then the new value function is

$$V_1(x,a) = \max_{x'} [\tilde{V}_1(x,x',a)].$$

Value Function Iteration and Constraints

- Everything is the same as before, except we need to adjust the value function to prohibit choices that violate the constraint.
- Just use an if statement.
- If *constraint violated* then *new value function is really really negative (e.g. -100000)*.
- Agents will avoid these control choices since it leads to really negative utility. Easy!
- Obviously though, if you have an exogenous lower bound (say $\underline{a}$ for assets), then you can just choose the lower bound on your grid to be that (as we will in the problem set).

Roadmap

Setup

- We have all the ingredients to solve for a stationary equilibrium of one of these models.
- Huggett's (1993) model studies an endowment economy, while Aiyagari's (1994) studies a production economy.
- Today we'll talk about an endowment economy: I will mention a couple of features of Aiyagari (1994) though.
- We'll go into the production economy in a bit more detail in the next lecture.

Setup

- Assume a unit measure of agents.
- Endowment economy with the usual preferences over consumption

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t U(c_t)$$

Setup

- Idiosyncratic uncertainty with regard to employment status/earnings.
- Two states: employed (e) and unemployed (u).
- Denote a households' state for a given period by $s_t \in \{e, u\}$.
- Denote their earnings by $y_t(s_t)$.
- If $s_t = e$ then $y_t(e) = 1$.
- If $s_t = u$ then $y_t(u) = b < 1$.
- Assume a Markov transition process across the states

$$\pi(s'|s) = \text{Prob}(s_{t+1} = s'|s_t = s)$$

Setup

- The households can borrow and save through one period discount bonds.
- These bonds are non-state contingent.
- Households can borrow subject to exogenous limit $a_t \geqslant \underline{a}$.
- Discount bond price at time t given by $q_t < 1$.
- No aggregate uncertainty: means that q will be constant in the RCE.
- We'll solve for q endogenously such that there is equilibrium in the bond market.

Household Problem

- Household solves

$$\max_{\{c_t, a_{t+1}\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t U(c_t)$$

subject to

$$c_t + q_t a_{t+1} \leqslant a_t + y_t$$
$$a_{t+1} \geqslant \underline{a}$$

$\forall t, s_t$.

Natural Borrowing Limit

- While Huggett (1993) takes $\underline{a}$ as given, Aiyagari (1994) allows households to borrow up to their natural borrowing limit.
- This is defined as the most that can be borrowed and repaid with probability one across all possible histories.
- If the household is unemployed today, their budget constraint is

$$\begin{aligned} c_t + q_t a_{t+1} &= a_t + b \\ \Rightarrow c_t + q_t[c_{t+1} + q_{t+1}a_{t+2} - b] &= a_t + b \\ \Rightarrow c_t + q_t c_{t+1} + q_t q_{t+1}[c_{t+2} + q_{t+2}a_{t+3} - b] - q_t b &= a_t + b. \end{aligned}$$

Natural Borrowing Limit

- Using the fact that $q_t = q$ in the stationary equilibrium and continuing to iterate gives

$$\lim_{T\to\infty}\left[\sum_{\eta=0}^{T-1} q^{\eta} c_{t+\eta} - q^{\eta} b + q^{T} a_{t+T}\right] = a_t$$

- To rule-out Ponzi schemes, we require that

$$\lim_{T\to\infty} q^{T} a_{t+T} \geqslant 0.$$

- Therefore the most the household could repay is

$$\begin{aligned}\underline{a}^{N} &= -\sum_{\eta=0}^{\infty} q^{\eta} b \\ &= -\frac{b}{1-q}\end{aligned}$$

Natural Borrowing Limit

- Notice that q is an endogenous object though!
- So this constraint can be a pain in the neck.
- Recall earlier that we saw $R\beta < 1$ where R is a gross return on bonds.
- Notice that the price is related to return via $q = \frac{1}{R}$.
- Therefore $q > \beta$.
- We can then instead use $\underline{a}^{\beta} = -\frac{b}{1-\beta}$ as the limit.
- For the rest of this lecture, we'll go back to taking the limit as exogenously given by $\underline{a}$.

Household's Recursive Problem

- We can write the Bellman equation as

$$v(s, a) = \max_{a'} U(y(s) + a - qa') + \beta\mathbb{E}_{s'}[v(s', a')]$$

where $a' \in \left[\underline{a}, \frac{y(s)+a}{q}\right]$.

- Note that the upper-limit is to ensure the period utility function is properly defined.

- The solution to the problem will be a policy function for the asset control $a' = g(s, a)$.

Cross-Sectional Distribution

- Everything's been pretty standard so far.
- I.e. recursive household problem: thus far all we've done is add-in an occasionally binding borrowing constraint.
- But notice that households in the population will generally differ in their asset holdings, since their states will generally differ.
- Therefore we need to track what the cross-sectional dispersion is over the state space.

Cross-Sectional Distribution

- Denote the measure over assets and employment status at time t by $\mu_t(\mathscr{S}, \mathscr{A})$ where $\mathscr{S}$ and $\mathscr{A}$ are the spaces for employment and assets respectively.
- Notice that I used the word measure.
- In a general RCE, this measure need not integrate to unity.
- E.g. if there is entry and exit of agents into the model: its integral may sum to a number greater than one.
- In this instance though, it'll integrate to one since we assumed a unit mass of agents: we can interpret μ_t as a probability distribution.

Cross-Sectional Distribution

- Given an arbitrary starting point, we can utilise the policy function for assets and the employment transition probabilities to map the evolution of the cross-section.
- In the steady state, $\mu_t \to \mu$.
- But in general, it will evolve as

$$\mu_{t+1}(\mathscr{S}, \mathscr{A}) = \int_{\mathscr{S}, \mathscr{A}} \left\{ \int_{s,a} \mathbb{1}_{a'=g(s,a)}(s, a)\pi(s'|s)\mu_t(ds, da) \right\} ds' da' \quad (1)$$

Cross-Sectional Distribution

- I've always found this measure-theoretic notation very confusing.
- It helps to think about how to find this object on a computer with a discretised state space.
- Say that $\mathscr{S} = \{e, u\}$ and $\mathscr{A} = \{a^1, a^2, ..., a^N\}$ are how we've discretised the spaces.
- You can follow the following procedure to implement this on the computer (I'll list computer variables in italics).

Cross-Sectional Distribution

(1) Take whatever staring point for μ_t that you like, (e.g. uniform across the grids: $\mu_t = \frac{1}{2 \times N}$ in our example). Call this vector *mu* in your code. It's of size (2,N).

Begin while loop (while *criterion* $> \epsilon$)

(2) Create a new vector called *mup* (for mu prime). This is also a vector of size (2,N). Initialise it equal to zero.

Begin do loops for each variable (loop over s and a).

(3) Update *mup* as follows

$$mup(s', g(s, a)) = mup(s', g(s, a)) + \pi(s'|s) \times mu(s, a)$$

End do loops for each variable (end s and a loops).

(4) Update *criterion* as

$$criterion = \sup || mu \; - \; mup||$$

(5) Update *mu* to take the *mup* (new measure)

$$mu = mup$$

End while loop

Cross-Sectional Distribution

- Assuming that your problem is well-behaved, the difference between *mu* and *mup* should converge to something close to zero (less than ϵ small).
- The resulting *mu* vector will be your steady state distribution in this case.

Recursive Competitive Equilibrium Definition

- A steady state recursive competitive equilibrium is a set of policy functions (c, a'), a price q and an invariant cross-section μ such that
  - For the given q, the recursive program solves the household's optimisation problem,
  - Given the stationary cross-section, goods and asset markets clear

$$\int_{\mathcal{S},\mathcal{A}} [c(s,a) - y(s)] d\mu = 0$$

$$\int_{\mathcal{S},\mathcal{A}} [a'(s,a)] d\mu = 0,$$

  - μ is a stationary measure: meaning that plugging μ into the right-side of (1) will return μ on the left-side.

Algorithm for Finding RCE

- Excess demand functions are your best friend!
- Most basic way to solve this:
  - **(a)** Conjecture a price q_0,
  - **(b)** Given q_0, solve the household's recursive problem. Gives you the corresponding policy functions for controls,
  - **(c)** Iterate on the cross-sectional law of motion (1) until it converges,
  - **(d)** Find the excess demand for assets. If it's close to zero, stop. If not, update the bond price as follows:
    - If excess demand is positive, increase q,
    - If excess demand is negative, decrease q.
  - **(e)** Return to step (b) using your updated price.

Algorithm for Finding RCE

- In the updating step, (d), you can just use the bisection method!
- E.g. set $q_{ub} = 1$ and $q_{lb} = \epsilon$ and then the initial guess is $q_0 = \frac{q_{ub}+q_{lb}}{2}$.
- If your excess demand is positive, set $q_{lb} = q_0$.
- If your excess demand is negative, set $q_{ub} = q_0$.

Quick Aside on Programming Languages

- The more prices you have as endogenous in a problem like this, the more loops you have.
- The number of iterations you need to undertake grows exponentially.
- Fortran is the king of loopy computing.
- If you have a complicated problem with many endogenous things, Fortran on a high-performance supercomputer can crush it much faster than Matlab.

Roadmap

Motivation

- We now know how to find the steady state RCE for a heterogeneous agent model.
- How can we think about transitions in this context?
- Same spirit as the representative agent model, but now we also have the cross-sectional measure and excess demand to deal with...

Implementation

- This is again a version of the shooting algorithm.
- Say that a parameter changes once and for all: here let's think about an increase in b [unemployed income] to $\tilde{b}$.
- I.e. the unemployed start earning more.
- We can solve easily for the new steady state using the algorithm in the previous section.
- What happens along the transition though?

Implementation

- What's the key equilibrium object here?
- The price of bonds q_t: it determines our agents' savings behaviour.

Shooting Algorithm with Heterogeneity

(I) Solve for the initial steady state RCE (corresponding to b): gives you q as well as $v(s, a)$.

(II) Solve for the new steady state RCE (corresponding to $\tilde{b}$): gives you $\tilde{q}$ as well as $\tilde{v}(s, a)$.

(III) Fix a number of time periods to convergence — T. This is something that you choose.

(IV) Guess a sequence of bond prices $\{\hat{q}_t\}_{t=1}^{T}$ such that $\hat{q}_1 = q$ and $\hat{q}_T = \tilde{q}$.

Shooting Algorithm with Heterogeneity

(V) Shoot backwards from time $t = T$ given the terminal value function for the households, $\tilde{v}(s, a)$. That is: solve the household's optimisation problem for $T-1, T-2, \ldots, 1$ working backwards from the final period. The value of q_t to use in each optimisation problem is that, which we conjectured — $\hat{q}_t$.

(VI) Shoot forwards using the law of motion for the cross-sectional measure (1). This gives you the evolution of the cross-section over time.

(VII) Check to see if the market for asset holdings clears for every time period along the transition path. That is: check if

$$\max_{1 \leqslant t \leqslant T} |ED_t(\hat{q}_t)| = \max_{1 \leqslant t \leqslant T} \left| \int_{\mathcal{S},\mathscr{A}} a'_t(s, a) d\mu_t \right| < \epsilon$$

where a'_t here denotes the policy function for a' at time t.

Shooting Algorithm with Heterogeneity

(VIII) If the asset market doesn't clear for every t, then update the price sequence. This is a little more art than science. Usually you'll use something to the effect of

$$\hat{q}_t^{new} = \hat{q}_t + \eta ED_t(\hat{q}_t)$$

where η is some positive constant. That is: if people save too much, drop the interest rate, meaning increase the bond price and vice-versa.

(IX) Go back to step (V) using $\{\hat{q}_t^{new}\}_{t=1}^{T}$ as your new bond price sequence.

Welfare Effects of Policy Changes

- In this example, I've applied the shooting algorithm to a change in a parameter.
- Usually it's most important for considering policy changes though. Was a particular policy change a good or bad thing from the households' point of view?
- Say that the change to $\tilde{b}$ came about due to some government intervention.
- How much better-off does it make households relative to the initial stead state?

Welfare Effects of Policy Changes

- Use consumption equivalent variation.
- Conditional: given an initial state (s, a), how much more consumption must we give a household in the initial steady state to give him just as much utility as he gets along the transition and new steady state onwards?
- This gives a real measure: something that makes sense to compare quantitatively.
- Comparing utils makes no sense: need to convert it to something cardinal.

Welfare Effects of Policy Changes

- Say that $\{\tilde{c}_t\}_{t=1}^{\infty}$ denotes the level of consumption post policy change for a consumer with starting state (s, a).

- Compute the consumption equivalent measure, $\omega(s, a)$ as

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u([1 + \omega(s, a)] c_t^*) = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(\tilde{c}_t).$$

  where c_t^* is their pre policy change level of consumption.

- This approach takes account of welfare differences — both between steady states and the transition.

Roadmap

Takeaways

- If you can implement all these techniques from today, you open-up a world of truly exciting research questions that you can answer.
- Get to work!