The Standard Incomplete Markets Model

Manuel Macera | 2026

Where we are

- Last lecture: a single household self-insures with a risk-free bond. If $\beta < q$, its assets stay bounded and have a long-run distribution.
- With a continuum of such households, that distribution is the **cross section**, and its mean is the **supply of savings**, $A(r)$.
- Today we close the model: the interest rate adjusts until $A(r)$ meets the demand for savings from outside the household sector. That demand is zero in an exchange economy and equals firms' demand for capital in a production economy.
- Two things are new:
  - the equilibrium has **no closed form**, so we must compute it;
  - the model makes predictions about the **wealth distribution**, which we can confront with data.

Plan

1. **Equilibrium**: an exchange economy *(Huggett, 1993)* and a production economy *(Aiyagari, 1994)*.

2. **Computation**: a search over the interest rate and, for each rate, two fixed-point problems: the household's policy and the distribution.

3. **Predictions**: what the two models say about interest rates, saving, the wealth distribution and marginal propensities to consume.

4. **An extension**: entrepreneurs, collateral constraints and the top of the wealth distribution.

Part 1: equilibrium

A continuum of households

- A unit mass of households, indexed by $i \in [0, 1]$, each solving the income fluctuation problem.
- Income follows the same Markov chain for everyone, $y \in \{y_1, \ldots, y_{N_y}\}$ with transition matrix $\Pi = [\pi_{ij}]$, $\pi_{ij} = \Pr[y' = y_j \mid y = y_i]$, but draws are **independent across households**.
- Let the column vector $\pi^\star$ be the **stationary distribution** of the chain, $\pi^\star = \Pi^\top \pi^\star$: $\pi_i^\star$ is the long-run probability of y_i.
- All shocks are idiosyncratic: there is **no aggregate risk**.
- A **law of large numbers** applies *(Uhlig, 1996)*: cross-sectional averages equal expectations, so
  - the share of households with income y_i is $\pi_i^\star$;
  - aggregates are deterministic, even though every household faces risk.

The household's problem in recursive form

- With incomplete markets, aggregation fails: aggregate saving depends on how wealth is distributed, not only on its total. So the distribution affects **prices**.
- Let λ be the distribution of households over individual states (a, y). The household's problem is

$$V(a, y; \lambda) = \max_{c,\, a'} \; u(c) + \beta\, \mathbb{E}[V(a', y'; \lambda') \mid y] \quad \text{s.t.} \quad c + q(\lambda)\, a' = y + a, \quad a' \geq \underline{a}.$$

- For the problem on the right-hand side to be well defined, the household must know everything it needs to solve it:
  - its **individual state**, (a, y);
  - today's λ, which sets today's price $q(\lambda)$;
  - tomorrow's λ', to evaluate the continuation value.
- How does the household know λ'?

The law of motion of the distribution

- We assume the household knows a law of motion, T, and add it to the problem:

$$
\begin{aligned}
V(a, y; \lambda) = \max_{c,\, a'} \;\; & u(c) + \beta\, \mathbb{E}\big[V(a', y'; \lambda') \mid y\big] \\
\text{subject to} \;\; & c + q(\lambda)\, a' = y + a, \quad a' \geq \underline{a}, \\
& \lambda' = T(\lambda).
\end{aligned}
$$

- Without T, the problem is not well defined: the household could not evaluate its continuation value.

- A demanding assumption: the household must forecast how a whole distribution moves.

The distribution is a state variable

- λ is an aggregate state: common to all households, and part of every household's state. It is infinite-dimensional, and in principle all of it matters.

- So the well-defined problem is **hard to solve**. The household's state includes a whole distribution, and T must be consistent with the decisions of all households, which themselves depend on T.

- We therefore start with **stationary** equilibria, where $\lambda^\star = T(\lambda^\star)$. The distribution is still a state, but it does not move, so the household pays attention only to prices.

How the distribution moves: on a grid

- Suppose assets live on a grid $\{a_1, \ldots, a_{N_a}\}$ and, for now, that the policy $a^\star(a_n, y_i)$ lands on a grid point. A state is a pair (a_n, y_i).
- The **transition matrix** Q gives the share of households at (a_n, y_i) that move to (a_m, y_j) next period:

$$Q\big((a_n, y_i), (a_m, y_j)\big) = \underbrace{\pi_{ij}}_{\text{from } y_i \text{ to } y_j} \quad \underbrace{\mathbf{1}\{a^\star(a_n, y_i) = a_m\}}_{\text{does the policy lead to } a_m?}$$

- Next period's distribution adds over today's states:

$$\lambda'(a_m, y_j) = \sum_{n,i} Q\big((a_n, y_i), (a_m, y_j)\big)\, \lambda(a_n, y_i), \qquad \text{i.e.} \qquad \lambda' = Q^\top \lambda.$$

- So T is a matrix multiplication.

How the distribution moves: in general

- With continuous assets, replace grid points by sets and sums by integrals.
- Q is a **transition function**: the share of households at (a, y_i) that move into a set $\mathcal{A} \times \mathcal{Y}$ next period,

$$Q((a, y_i), \mathcal{A} \times \mathcal{Y}) = \sum_{j:\, y_j \in \mathcal{Y}} \underbrace{\pi_{ij}}_{\text{from } y_i \text{ to } y_j} \quad \underbrace{\mathbf{1}\{a^\star(a, y_i) \in \mathcal{A}\}}_{\text{does the policy lead into } \mathcal{A}?}$$

- The law of motion is

$$\lambda'(\mathcal{S}) = T(\lambda)(\mathcal{S}) = \int Q(s, \mathcal{S})\, \lambda(\mathrm{d}s) \qquad \text{for every set } \mathcal{S}.$$

- On a computer, everything is discrete: the previous slide *is* what we compute.

Equilibrium in an exchange economy

- The only asset is a bond issued and held by households: it is in **zero net supply** *(Huggett, 1993)*.
- A household can borrow only if another one lends, so $\underline{a} < 0$ is needed for trade.

Stationary recursive competitive equilibrium

A value function V, policies $c^\star(a, y)$ and $a^\star(a, y)$, a distribution $\lambda^\star$ and a bond price q such that:

1. given q, V solves the Bellman equation and $c^\star, a^\star$ are its policies;
2. the bond market clears: $\int a^\star(a, y)\, \lambda^\star(\mathrm{d}a, \mathrm{d}y) = 0$;
3. the distribution is stationary: $\lambda^\star = T(\lambda^\star)$.

Huggett's question

- Huggett (1993) asks: why is the average real risk-free rate so **low**?
  - over 1889–1978 it averaged 0.8% a year *(Mehra and Prescott, 1985)*, while calibrated representative-agent economies predict a rate that is too high.
- His conjecture: **market incompleteness** matters. Households cannot insure idiosyncratic risk and can borrow only up to a limit.

- The exchange economy isolates this mechanism. Questions we will answer once we can compute the equilibrium:
  - how far below the rate of time preference does the interest rate fall?
  - how does it depend on the borrowing limit?
  - and on risk aversion?

Adding production

- *Aiyagari (1994)* adds production. Households supply θ efficiency units of labor at wage w, so $y = w\,\theta$, with θ Markov.
- A representative firm rents capital at rate R and labor at wage w, with $F(K, L)$ constant returns and depreciation δ:

$$R = F_K(K, L), \qquad w = F_L(K, L).$$

- A risk-neutral intermediary sells bonds of face value A, collects qA and rents it to the firm. Zero profits require

$$K = q\,A, \qquad (1 + R - \delta)\,q = 1 \iff R - \delta = r.$$

- Saving now has a use: assets must be **positive** in the aggregate, and $\underline{a} = 0$ is allowed.

Equilibrium with capital

Stationary recursive competitive equilibrium

A value function V, policies $c^\star(a,\theta)$ and $a^\star(a,\theta)$, prices (r,w) and a distribution $\lambda^\star$ such that:

1. given (r,w), V solves the Bellman equation and $c^\star, a^\star$ are its policies, with $y = w\,\theta$ and $q = 1/(1+r)$;
2. aggregates add up individual decisions: $K = q\int a^\star \,\mathrm{d}\lambda^\star$ and $L = \int \theta \,\mathrm{d}\lambda^\star$;
3. prices are marginal products: $r = F_K(K,L) - \delta$ and $w = F_L(K,L)$;
4. the distribution is stationary: $\lambda^\star = T(\lambda^\star)$.

Aiyagari's question

- Aiyagari (1994) asks: how important is uninsured individual risk for **aggregate saving**?
- It was an open question where capital comes from: life-cycle saving, bequests, or precaution. Several authors conjectured that precautionary saving could account for a large share of the US capital stock *(Zeldes, 1989)*.
- Why capital matters for this question: in an exchange economy, saving cannot change aggregate resources. With capital, precautionary saving builds capital and, through diminishing returns, lowers the interest rate.
- Questions we will answer once we can compute the equilibrium:
  - how much higher are capital and the saving rate than with complete markets?
  - how does the answer depend on the size and persistence of earnings risk, and on risk aversion?
  - how unequal is wealth, compared with earnings?

Equilibrium is one equation in one unknown

- Fix r, and let λ_r be the stationary distribution implied by the household's policy at that r. Then r pins down:
  - the **supply of savings**, $A(r) = \int a \, \mathrm{d}\lambda_r$;
  - the wage $w(r)$ and the **demand for capital** $K(r)$, from the firm's condition.
- Equilibrium is a zero of excess supply:

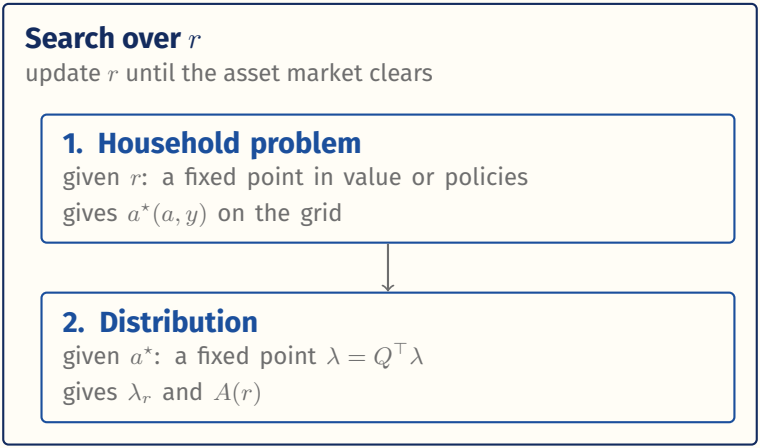
$$\text{Huggett: } A(r) = 0, \qquad \text{Aiyagari: } \frac{1}{1+r} A(r) = K(r).$$

- Existence: $A(r)$ is continuous, sits near $\underline{a}$ for low r, and diverges as $r \uparrow \beta^{-1} - 1$ (last lecture). In Aiyagari, $K(r)$ decreases in r. The curves cross.

- **Uniqueness is not guaranteed**, though in practice we find a single crossing.

Part 2: computing a stationary equilibrium

How to compute a stationary equilibrium



Steps 1 and 2 are fixed-point problems, solved once for every guess of r.

Preliminaries: the asset grid

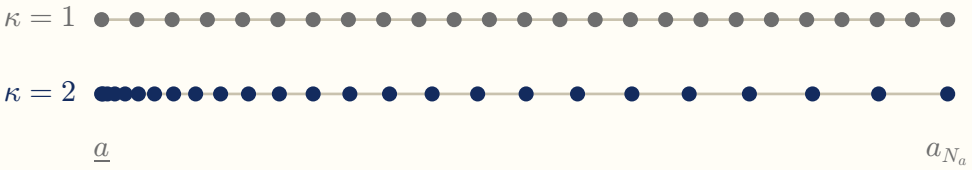
- A computer cannot handle a continuum of asset levels: we must **discretize** the domain of assets, $[\underline{a}, \infty)$, into a finite grid $\{a_1, \ldots, a_{N_a}\}$.

- A numerical solution is then a set of **matrices**: $a^\star(a_n, y_i)$, $c^\star(a_n, y_i)$ and $\lambda(a_n, y_i)$, of size $N_a \times N_y$.

- The grid starts at $a_1 = \underline{a}$ and ends at an **upper bound** a_{N_a} we must choose:
  - if $\beta < q$, assets are bounded, but we do not know the bound;
  - we pick a value and **check it** after solving (below).

Preliminaries: spacing the grid

- A convenient family: $a_n = \underline{a} + (a_{N_a} - \underline{a})\, x_n^{\kappa}$, with x_n evenly spaced on $[0, 1]$.
- With $\kappa = 1$ the grid is evenly spaced; with $\kappa > 1$ it puts more points near the borrowing limit.



- Ideally, we want more grid points where the functions we are solving for have more **curvature**.

Preliminaries: discretizing income

- Income processes estimated from data have a **continuous support**. We assume an AR(1) in logs as our leading example:

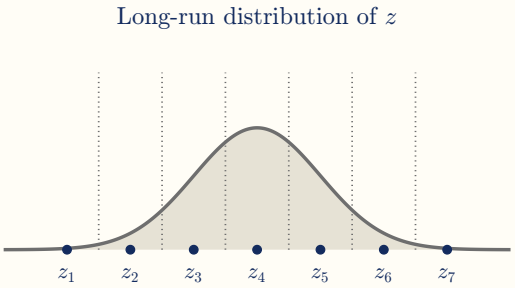
$$z' = \rho z + \varepsilon', \qquad \varepsilon' \sim N(0, \sigma_\varepsilon^2), \qquad y = \exp(z).$$

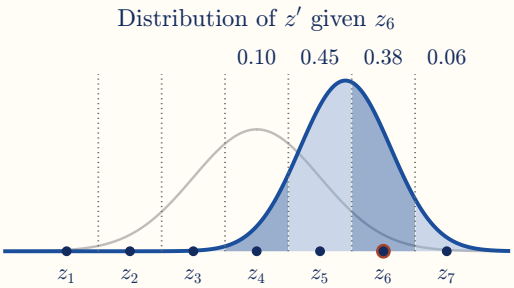
- We approximate it with a **Markov chain**, which means choosing two things:
  - a grid of points $\{z_1, \ldots, z_N\}$;
  - a transition matrix $\Pi = [\pi_{ij}]$.
- We want the chain to mimic the **stochastic properties** of the process.

- Two standard methods: Tauchen's and Rouwenhorst's *(Tauchen, 1986; Rouwenhorst, 1995)*.

Preliminaries: Tauchen's method

- **Grid**: N evenly spaced points on $\pm m$ unconditional standard deviations, $\sigma_z = \sigma_\varepsilon / \sqrt{1-\rho^2}$. Each point owns the cell around it.

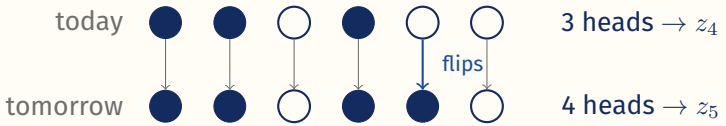




- **Transition matrix**: row i is the conditional density of z' given z_i, $N(\rho z_i, \sigma_\varepsilon^2)$, integrated over each cell. The tails go to z_1 and z_N.

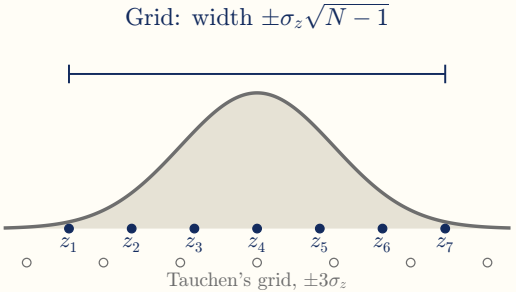
Preliminaries: Rouwenhorst's method, the idea

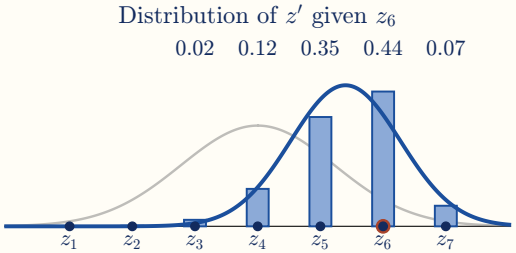
- Take the grid as given for a moment: N evenly spaced points.
- Take $N-1$ independent coins. Each period, each coin keeps its face with probability p and flips with probability $1-p$. The state is the number of heads, $k \in \{0, \ldots, N-1\}$: N values, one for each grid point z_{k+1}.



- Given N, we choose two things to match the AR(1) exactly: p for the **persistence**, and the width of the grid for the **variance**.
  - each coin, and so their sum, has autocorrelation $2p-1$: set $p = (1+\rho)/2$;
  - in the long run each coin is heads half the time, so the number of heads is Binomial$(N-1, \frac{1}{2})$, and a grid on $\pm\sigma_z\sqrt{N-1}$ gives variance σ_z^2 *(Kopecky and Suen, 2010)*.

Preliminaries: Rouwenhorst's method, grid and transition matrix





- Left: the grid, evenly spaced on $\pm\sigma_z\sqrt{N-1}$. For $N = 7$ it is narrower than Tauchen's $\pm 3\sigma_z$.
- Right: row 6 of Π. From k heads, tomorrow's heads are the heads that stay, Binomial(k, p), plus the tails that flip, Binomial$(N-1-k, 1-p)$. The curve is the AR(1)'s conditional density, scaled to the grid step.
- Instead of the normal areas that Tauchen's method uses, the chain matches the **persistence** and the **variance** exactly.

Output: which method to use

Persistence and standard deviation implied by each chain (target: $\sigma_z = 0.3$)

		Tauchen ($m = 3$)		Rouwenhorst	
True ρ	N	ρ	σ_z	ρ	σ_z
0.50	5	0.499	0.335	0.500	0.300
0.50	9	0.499	0.308	0.500	0.300
0.90	5	0.932	0.381	0.900	0.300
0.90	9	0.898	0.331	0.900	0.300
0.98	5	1.000	0.404	0.980	0.300
0.98	9	0.989	0.374	0.980	0.300

- Tauchen overstates risk, by more as persistence rises. For $\rho \geq 0.9$, the range that earnings estimates deliver, use **Rouwenhorst**.

A test economy for the household problem

- To compare algorithms we fix the price and solve one household problem:
  - CRRA utility with $\sigma = 2$, $\beta = 0.96$, $r = 3\%$ (so $\beta < q$), and $\underline{a} = 0$;
  - log income AR(1) with $\rho = 0.9$ and $\sigma_z = 0.3$, Rouwenhorst with $N_y = 7$.
- We measure accuracy with **Euler equation errors** *(Judd, 1992)*: off the grid, at points where the limit is slack,

$$\mathrm{EE}(a, y) = \log_{10} \left| 1 - \frac{u'^{-1}\left(\frac{\beta}{q}\, \mathbb{E}\left[u'(c^\star(a', y')) \mid y\right]\right)}{c^\star(a, y)} \right|.$$

- $\mathrm{EE} = -3$ means a mistake of \$1 for every \$1,000 spent.

Method 1: value function iteration

Algorithm: value function iteration on the grid

1. Guess $V_0(a_n, y_i)$.
2. For every (a_n, y_i), choose a' on the grid:

$$V_{k+1}(a_n, y_i) = \max_{a_m \geq \underline{a}} \; u(y_i + a_n - q\, a_m) + \beta \sum_j \pi_{ij}\, V_k(a_m, y_j).$$

3. Stop when $\|V_{k+1} - V_k\| < \varepsilon$; otherwise, go back to step 2.

- Robust: it needs neither concavity nor differentiability, and it always converges, since the Bellman operator is a contraction.
- But it is **slow and coarse**: the error shrinks by a factor β per iteration; each iteration evaluates N_a choices at each of $N_a N_y$ states, about $N_a^2 N_y$ operations; and the policy is a step function.

Speeding up VFI: Howard's improvement

- The expensive step is the maximization. But the **policy** converges long before the value function does.
- Howard's improvement: after each maximization, keep the policy $a^\star(a_n, y_i)$ fixed and update V H times,

$$V(a_n, y_i) \;\leftarrow\; u(y_i + a_n - q\, a^\star(a_n, y_i)) + \beta \sum_j \pi_{ij}\, V(a^\star(a_n, y_i), y_j),$$

  which needs no maximization.
- With $H = 50$, the number of maximizations falls from 443 to 15, for the same answer.
- But Howard does not improve accuracy: the policy still lives on the grid.

Value function iteration with a continuous choice

- Let a' take any value in $[\underline{a}, (y_i + a_n)/q)$, not only grid points. Two new ingredients:
  - **interpolate** the continuation value $\mathbb{E}[V_k(a', y') \mid y_i]$ between grid points;
  - **maximize** at every state with a one-dimensional optimizer, e.g. golden-section search.
- What it buys: the policy is no longer a staircase, so accuracy is no longer tied to the grid spacing.
- What it costs: $N_a N_y$ maximizations per iteration, each with dozens of evaluations of an interpolated function. Howard steps help here too.
- **Golden-section search assumes a single peak**: it needs a concave problem.

Method 2: policy function iteration

- Iterate on the Euler equation, not the Bellman equation.

Algorithm: policy function iteration

1. Guess a consumption policy $c_0(a_n, y_i)$.
2. For every (a_n, y_i), find the a', any value, that solves

$$u'(y_i + a_n - q\,a') = \frac{\beta}{q}\sum_j \pi_{ij}\,u'(c_k(a', y_j)),$$

   interpolating c_k between grid points. If the left side exceeds the right at $a' = \underline{a}$, set $a' = \underline{a}$.
3. Set $c_{k+1} = y_i + a_n - q\,a'$ and iterate until $\|c_{k+1} - c_k\| < \varepsilon$.

- Accurate, but each iteration solves $N_a N_y$ nonlinear equations: a root finder inside a loop.

Method 3: the endogenous grid method

- The idea of Carroll (2006): fix **tomorrow's** assets on the grid and solve for **today's**.
- If $a' = a_m$ is a grid point, the right side of the Euler equation needs no interpolation, and the Euler equation **can be inverted in closed form**:

$$\tilde{c}(a_m, y_i) = u'^{-1}\Big(\tfrac{\beta}{q} \textstyle\sum_j \pi_{ij}\, u'(c_k(a_m, y_j))\Big).$$

- The budget constraint then tells us **which assets today lead to that choice**:

$$\tilde{a}(a_m, y_i) = \tilde{c}(a_m, y_i) + q\, a_m - y_i.$$

- The points $\tilde{a}$ form an **endogenous grid**: no root finding and no maximization, only function evaluations.

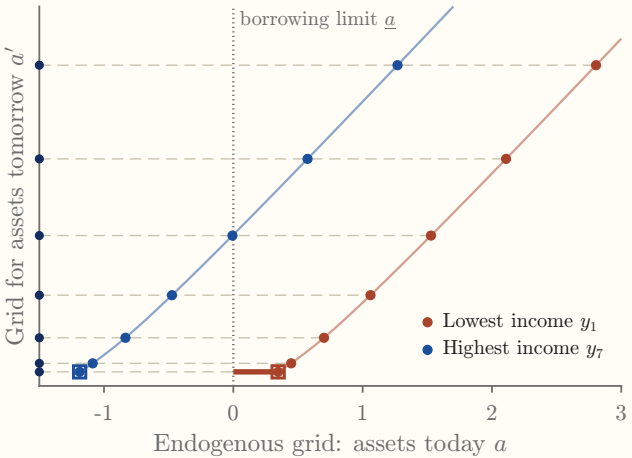
The endogenous grid method: the algorithm

Algorithm: endogenous grid method

1. Guess $c_0(a_n, y_i)$, e.g. $c_0 = y_i + a_n - q\,\underline{a}$.
2. For every (a_m, y_i), compute $\tilde{c}(a_m, y_i)$ and $\tilde{a}(a_m, y_i)$.
3. For every y_i, interpolate the pairs $(\tilde{a}(a_m, y_i),\ a_m)$ to get $a'(a_n, y_i)$ on the original grid.
4. Where $a_n < \tilde{a}(a_1, y_i)$, below the lowest point of the endogenous grid, the limit binds: set $a'(a_n, y_i) = \underline{a}$.
5. Set $c_{k+1} = y_i + a_n - q\,a'$ and iterate until $\|c_{k+1} - c_k\| < \varepsilon$.

- Why step 4: at $\tilde{a}(a_1, y_i)$ the Euler equation holds exactly with $a' = \underline{a}$. With less wealth, the household would like to borrow more, but cannot.

Output: the endogenous grid



- Each grid point for a', on the vertical axis, maps into a point $\tilde{a}$ of the endogenous grid.
- The square marks $\tilde{a}(a_1, y_i)$. Between the borrowing limit and the square, the limit binds and $a' = \underline{a}$.
- With low income the limit binds up to $a = 0.35$; with high income it never binds.

Output: the race

Method	N_a	Iterations	Time (s)	Euler error, $\log_{10}$ Mean	Euler error, $\log_{10}$ 95th pct.
VFI	1,000	443	3.59	-2.4	-2.0
VFI + Howard	1,000	15	0.27	-2.4	-2.0
VFI, continuous	200	443	3.28	-6.1	-5.1
VFI, continuous + Howard	200	17	0.20	-6.1	-5.1
Policy function iteration	200	382	1.52	-6.9	-5.2
EGM	200	382	0.12	-6.9	-5.2
EGM	1,000	381	**0.18**	-8.3	-6.5

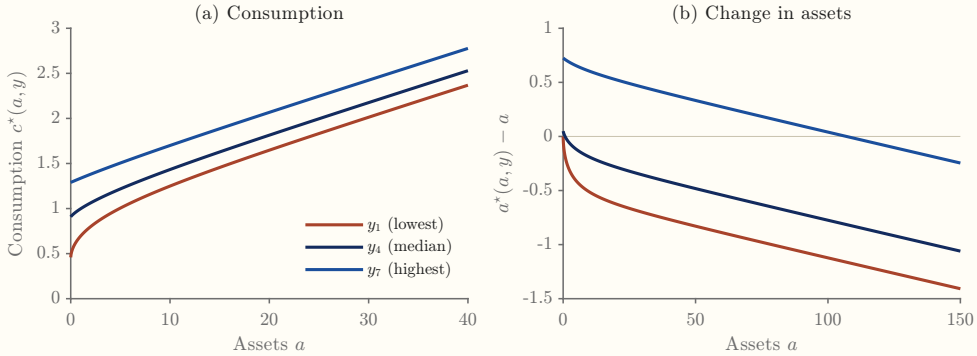
- With a continuous choice, VFI is accurate too. On the same grid, policy function iteration is faster: finding a zero is cheaper than maximizing.
- EGM gives the same answer as policy function iteration, 13 times faster. It is as fast as grid-based VFI with Howard, and **about a million times more accurate**.

MATLAB on a laptop; tolerance 10^{-8}; errors on 5,000 points off the grid.

When EGM is not enough

- EGM wins when the problem has one endogenous state, one continuous choice, a concave value function and an invertible marginal utility.
- Adding exogenous states, such as more income points, is cheap: the cost grows only linearly.
- It struggles with:
  - more than one endogenous state, e.g. two assets: the endogenous grid comes out as scattered points;
  - more than one continuous choice, e.g. consumption and hours;
  - non-concave problems, e.g. discrete choices such as default or occupation *(Iskhakov et al., 2017)*;
  - no closed-form inverse of marginal utility: we are back to a root finder.
- Then we fall back on value function iteration, on the grid if the problem is not concave, or on policy function iteration, which is still accurate.

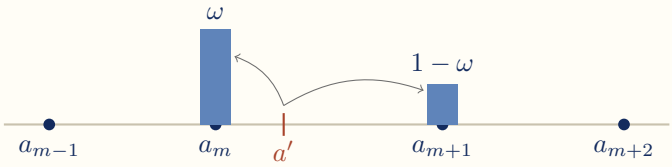
Output: policies, and a check on the grid



- Assets fall above some level for **every** income state, so the upper bound is not binding.
- The highest-income household keeps saving up to $a \approx 107$: an upper bound of $a_{N_a} = 40$ would have **distorted its policy**.

The distribution: two ways to compute it

- **Monte Carlo**: simulate many households for many periods with the policy, and use the cross section at the end.
  - simple, but noisy: the error falls only with the square root of the number of households.
- **Histogram**: track the mass at every grid point. A household whose choice a' falls between two grid points is split between them by a lottery:



$$\omega = \frac{a_{m+1} - a'}{a_{m+1} - a_m}$$

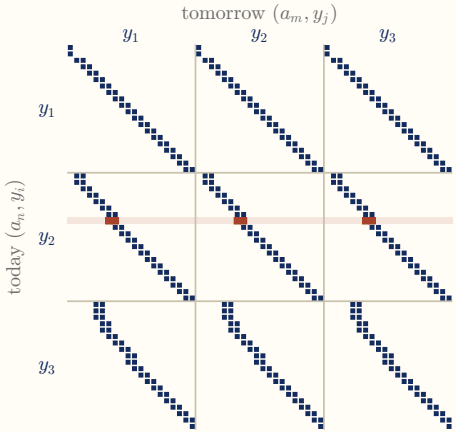
- The closer grid point gets more mass, and the mean is exact: $\omega\, a_m + (1-\omega)\, a_{m+1} = a'$. The histogram has no sampling noise.

The histogram method: the transition matrix

- Order the states as $s = (a_n, y_i)$ and fill an $N_a N_y \times N_a N_y$ **sparse** matrix:

$$Q((a_n, y_i), (a_m, y_j)) = \pi_{ij}\, \omega_{nim},$$

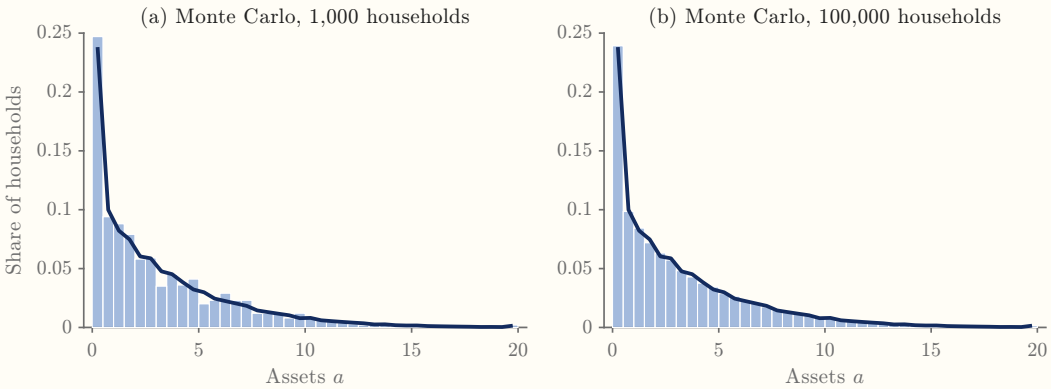
  where ω_{nim} is the lottery weight that $a^\star(a_n, y_i)$ puts on a_m.
- Each row has at most $2N_y$ nonzero entries, as in the highlighted row: two neighboring grid points for each income state tomorrow.



The histogram method: the stationary distribution

- The stationary distribution solves $\lambda = Q^\top \lambda$. Three ways to find it:
  - iterate $\lambda_{k+1} = Q^\top \lambda_k$ from any guess (827 iterations here);
  - compute the eigenvector of $Q^\top$ for the unit eigenvalue;
  - solve the linear system $(Q^\top - I)\lambda = 0$, with one equation replaced by $\sum_s \lambda_s = 1$.
- All three take a fraction of a second.

Output: histogram against Monte Carlo



- Line: histogram method, mean assets 3.114, in 0.1 seconds. With 1,000 households, Monte Carlo gives 3.072; with 100,000, 3.109 in 3 seconds.
- Noise in $A(r)$ is deadly in the search over r: it makes excess supply jagged in r.

The search over r: clearing the market

- Excess supply $A(r)$ (Huggett) or $qA(r) - K(r)$ (Aiyagari) is one function of one variable.
- **Bisection** on r is slow but safe. The bracket:
  - from above, $r < \beta^{-1} - 1$, where $A(r)$ diverges;
  - from below, $r > -\delta$ in Aiyagari. Under an ad hoc limit, the household problem is well defined for $r \leq 0$ too, so in Huggett the bracket can include negative rates.
- Alternatives in Aiyagari: iterate on K, updating $K_{k+1} = (1 - \xi)K_k + \xi\, qA(r(K_k))$ with damping ξ.
- **Warm starts**: start each household solution from the previous c. Late in the bisection the price barely moves, so the household problem converges in a few iterations.

The whole algorithm

Algorithm: stationary equilibrium

0. Build the asset grid and the income chain (y, Π). Set a bracket $[r_L, r_H]$.
1. Set $r = (r_L + r_H)/2$, $q = 1/(1+r)$, and $w(r)$ from the firm's condition.
2. Solve the household problem by EGM, starting from the last c: get $a^\star(a_n, y_i)$.
3. Build Q with lottery weights and solve $\lambda = Q^\top \lambda$.
4. Compute excess supply, $E(r) = q \sum_{n,i} a^\star(a_n, y_i)\, \lambda(a_n, y_i) - K(r)$.
5. If $|E(r)| < \varepsilon$, stop. If $E(r) > 0$, set $r_H = r$; otherwise set $r_L = r$. Go back to step 1.

- For Huggett, set $K(r) = 0$ and $w = 1$.

A checklist

- **Tolerances**: tighter for the household problem and the distribution than for the search over r. Otherwise, their numerical noise flips the sign of excess supply.
- **Top of the grid**: $a^\star(a, y) < a$ near a_{N_a} for every y, and no mass at a_{N_a}.
- **Accuracy**: Euler errors off the grid.
- **Market clearing**: report excess supply at the solution, and check Walras' law: the goods market clears too.
- **Replicate** a known result before trusting a new one. We do this next.

Part 3: Huggett (1993) and the risk-free rate

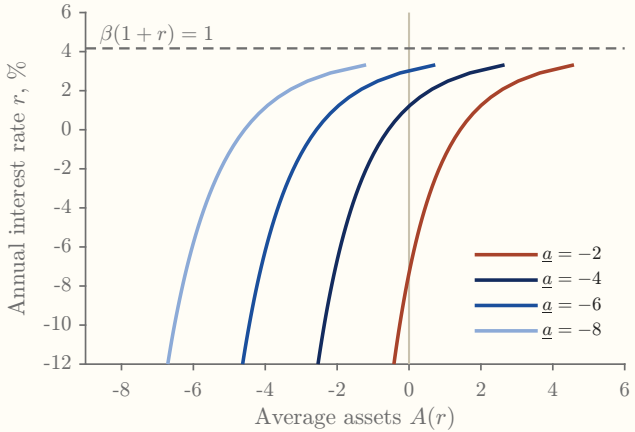
Calibration

- Back to Huggett's question: how much does self-insurance depress the risk-free rate?

- A model period is two months. Income is an employment state:
  - $y \in \{0.1, 1\}$: not employed and employed;
  - $\pi(1 \mid 1) = 0.925$ and $\pi(1 \mid 0.1) = 0.5$: a spell without work lasts two periods on average;
  - the standard deviation of annual earnings is 20% of the mean.
- $\beta = 0.99322$ (0.96 per year), CRRA with $\sigma = 1.5$, and limits $\underline{a} \in \{-2, -4, -6, -8\}$. Annual income is 5.3, so $\underline{a} = -5.3$ is one year of income.

- With complete markets, the annual rate is $\beta^{-6} - 1 = 4.2\%$.

Prediction: the supply of savings



- Each curve is $A(r)$ for one limit; equilibrium is where it crosses $A=0$.
- A looser limit shifts $A(r)$ **left**: households borrow more at every rate, so the rate must rise to clear the market.
- All curves bend up toward $\beta(1+r)=1$.

$\sigma=1.5$. *Rates are annual. With an ad hoc limit,* $A(r)$ *is well defined for* $r \leq 0$, *i.e.* $q \geq 1$.

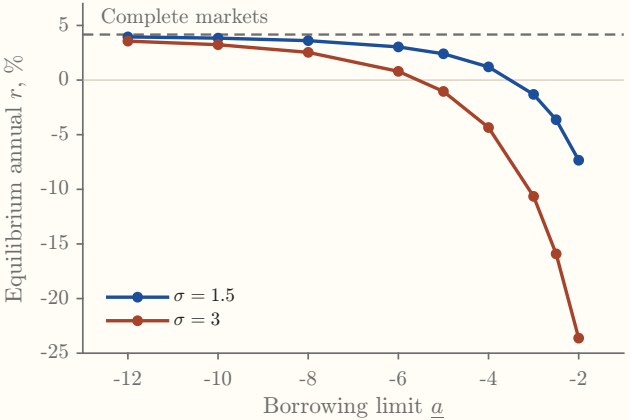
Replication: Huggett's Table 1

Equilibrium annual interest rate, %				
Borrowing limit $\underline{a}$	-2	-4	-6	-8
$\sigma = 1.5$				
Huggett (1993)	-7.1	2.3	3.4	4.0
This lecture	-7.3	1.2	3.0	3.6
$\sigma = 3$				
Huggett (1993)	-23.1	-2.7	1.8	3.7
This lecture	-23.6	-4.3	0.8	2.5

- Same ranking, same orders of magnitude. Our rates are lower, and the gap **survives a finer grid**: at Huggett's price for $\underline{a} = -4$, our excess supply is 0.87, far above his tolerance of 0.0025.
- A useful habit: when a replication misses, find out whether the gap comes from your code or from theirs.

Huggett's rates for $\sigma = 3$ are computed from his reported prices.

Prediction: tighter limits, lower rates



- As the limit loosens, the rate approaches the complete-markets rate, but stays **below** it.
- As it tightens, the rate falls fast: the less households can borrow, the lower the rate must be to stop them from saving.
- More risk aversion means more precaution, and lower rates.

All limits shown are ad hoc: the natural limit, $-y_{\min}/(1-q)$, lies below -15 at every rate shown.

What we learn from Huggett

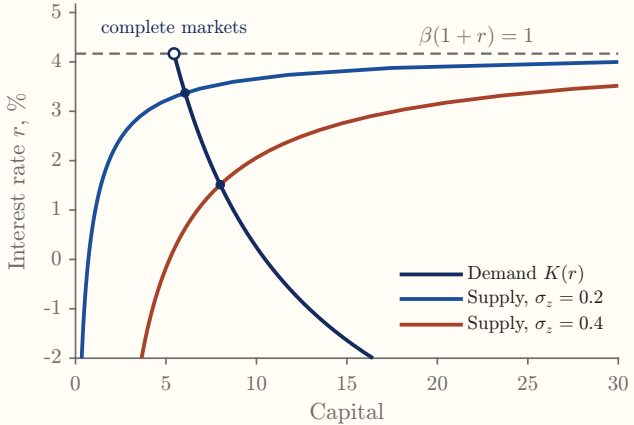
- Uninsurable risk lowers the risk-free rate below the rate of time preference, $\beta(1 + r^\star) < 1$.
- The size of the effect depends on how tightly households can borrow: with one year of income as a limit, the rate falls by about 1.5 points.
- A low risk-free rate can come from **market incompleteness**, without implausible preferences.
- Next: what changes when saving can build capital.

Part 4: Aiyagari (1994), saving and capital

Calibration

- Back to Aiyagari's question: how much does self-insurance raise aggregate saving?
- Annual model. Cobb–Douglas technology, $F(K, L) = K^{\alpha} L^{1-\alpha}$, with $\alpha = 0.36$ and $\delta = 0.08$; $\beta = 0.96$; $\underline{a} = 0$.
- Log labor efficiency is an AR(1) with persistence $\rho \in \{0, 0.3, 0.6, 0.9\}$ and standard deviation $\sigma_z \in \{0.2, 0.4\}$. Relative risk aversion $\sigma \in \{1, 3, 5\}$.
- Aiyagari uses Tauchen with seven states; so do we, to replicate him.
- With complete markets, $r = \beta^{-1} - 1 = 4.17\%$ and the saving rate is $s = \delta\alpha/(r + \delta) = 23.7\%$.

Prediction: savings and capital



Tauchen, $\rho = 0.9$, $\sigma = 3$, as in Aiyagari.

- Demand $K(r)$ decreases with r; supply $qA(r)$ increases and diverges as $\beta(1+r) \to 1$.
- With complete markets, supply would be flat at $\beta(1+r) = 1$: the open dot, $K = 5.45$.
- With risk, capital is **higher** and the rate **lower**: $K = 6.05$ and $r = 3.37\%$ ($\sigma_z = 0.2$); $K = 8.00$ and $r = 1.51\%$ ($\sigma_z = 0.4$).

Replication: Aiyagari's Table II

Equilibrium interest rate, % ($\sigma_z = 0.4$)

	This lecture			Aiyagari (1994)		
ρ	$\sigma = 1$	$\sigma = 3$	$\sigma = 5$	$\sigma = 1$	$\sigma = 3$	$\sigma = 5$
0.0	4.06	3.78	3.45	4.06	3.78	3.42
0.3	3.98	3.49	2.93	3.96	3.42	2.80
0.6	3.80	2.91	2.00	3.76	2.78	1.81
0.9	3.40	1.51	-0.09	3.31	1.29	-0.35

- Close everywhere, and **almost identical** when shocks are not persistent; the gaps grow with ρ and σ.
- Relative to 23.7% with complete markets, incomplete markets raise the saving rate by 0.2 points ($\rho = 0$, $\sigma = 1$) to 12.7 points ($\rho = 0.9$, $\sigma = 5$). Aiyagari: 0.2 to 14.0.

The panel with $\sigma_z = 0.2$ matches within 0.15 points.

Output: the discretization matters

Equilibrium interest rate, %

	$\sigma_z = 0.2$			$\sigma_z = 0.4$		
$\rho = 0.9$	$\sigma = 1$	$\sigma = 3$	$\sigma = 5$	$\sigma = 1$	$\sigma = 3$	$\sigma = 5$
Tauchen, 7 states	3.95	3.37	2.67	3.40	1.51	-0.09
Rouwenhorst, 7 states	4.01	3.58	3.05	3.59	2.12	0.73

- Tauchen's chain is riskier than the AR(1) it approximates (first table of Part 2), so it **overstates** precautionary saving.
- With persistent shocks and high risk aversion, the bias reaches **0.8 percentage points** of interest rate: as large as the effects we want to measure.
- From now on: Rouwenhorst, $\rho = 0.9$, $\sigma_z = 0.4$, $\sigma = 3$. Then $r = 2.12\%$ and $K/Y = 3.56$.
- The answers depend heavily on ρ. We discuss how persistent income really is when we study inequality.

How much more do households save?

- In a steady state, investment replaces depreciation, and the firm sets $F_K = r + \delta$. With Cobb–Douglas,

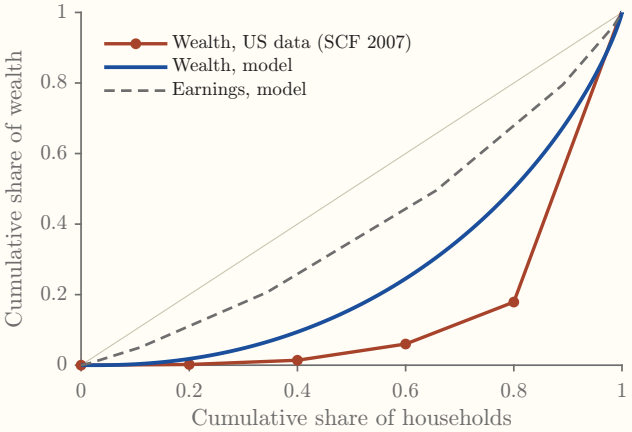
$$s = \frac{\delta K}{Y} = \frac{\delta \alpha}{r + \delta}.$$

- The formula is the same with complete and incomplete markets. Risk changes saving **only through** r:

$$s^{IM} - s^{CM} = \frac{\delta\alpha\left(\beta^{-1} - (1 + r^{IM})\right)}{(r^{IM} + \delta)(\beta^{-1} - 1 + \delta)} > 0.$$

- For moderate risk, the increase is at most about one point. It is large only with very persistent shocks and high risk aversion.
- The model's lasting influence comes from its other predictions: the **distribution of wealth**, which Aiyagari himself notes is less unequal than in the data, and, in later work, **marginal propensities to consume**. We turn to these next.

Prediction: the distribution of wealth



- The model turns an earnings Gini of 0.22 into a wealth Gini of **0.48**: precautionary saving amplifies inequality.
- But the data show a wealth Gini of 0.78.
- The gap is at the **top**: the richest 1% hold 4.7% of wealth in the model, and 33.5% in the data.

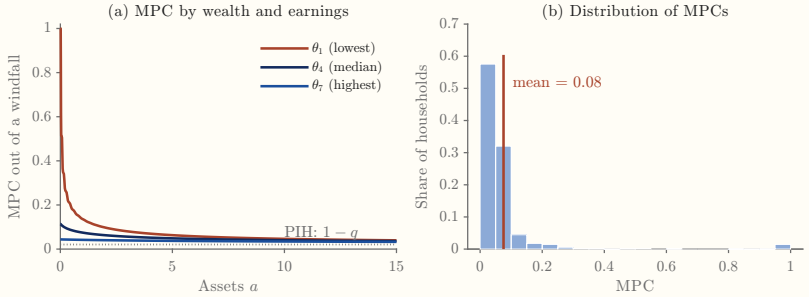
Data: Survey of Consumer Finances 2007, from Krueger, Mitman and Perri (2016).

Prediction: the wealth distribution in numbers

	Share of wealth held by quintile, %						
	Q1	Q2	Q3	Q4	Q5	Top 1%	Gini
US data (SCF 2007)	0.2	1.2	4.6	11.9	82.5	33.5	0.78
Aiyagari economy	1.8	7.6	15.1	25.7	49.8	4.7	0.48

- The bottom 40% hold almost nothing in the data; in the model, they hold 9%.
- The middle class is too rich, and the rich are **not rich enough**.
- Why? Households save to self-insure. Once they hold a buffer, the motive fades, and with $\beta(1+r) < 1$ they stop accumulating. Nothing in the model makes the rich keep saving.

Prediction: marginal propensities to consume



- A **windfall** is a one-time, unexpected transfer Δ added to today's assets. The annual MPC out of it is $[c^\star(a+\Delta, y) - c^\star(a, y)]/\Delta \approx \partial c^\star/\partial a$, with Δ equal to 1% of the wage.
- Households at the limit consume all of a windfall; the MPC then falls quickly toward the permanent income value.
- Few households are near the limit, so the mean annual MPC is only 0.08, well below the 0.15–0.25 per *quarter* found for tax rebates *(Kaplan and Violante, 2014)*.

What we learn from Aiyagari

- Self-insurance raises capital and lowers the interest rate; the effect on the saving rate is modest unless risk is large and persistent.
- The model generates a wealth distribution from first principles, and it is more unequal than earnings.
- Two of its predictions are far from the data:
  - less wealth at the top;
  - fewer households with little wealth, hence a lower average MPC.
- The second matters for fiscal and monetary policy, later in the course. Next, we extend the model and look again at the first.

Part 5: an extension with entrepreneurs

Why entrepreneurs

- In US data, families that own and run a business are about 14% of families. They earn 22% of income and hold 40% of wealth *(Quadrini, 2000)*.
- They are concentrated at the **top** of the wealth distribution.
- Running a firm gives rich households a reason to keep saving: if borrowing is limited, a large firm needs the owner's own wealth, and the return on it is high.
- We add entrepreneurs to the Aiyagari economy, keep everything else the same, and solve for the new equilibrium.

Entrepreneurs: the environment

- Each household has, besides θ, an entrepreneurial ability z, drawn from a Pareto distribution; each year, with some probability, it draws a new one.
- At the start of each period, the household chooses to work, for $w\,\theta$, or to run a firm, for profits

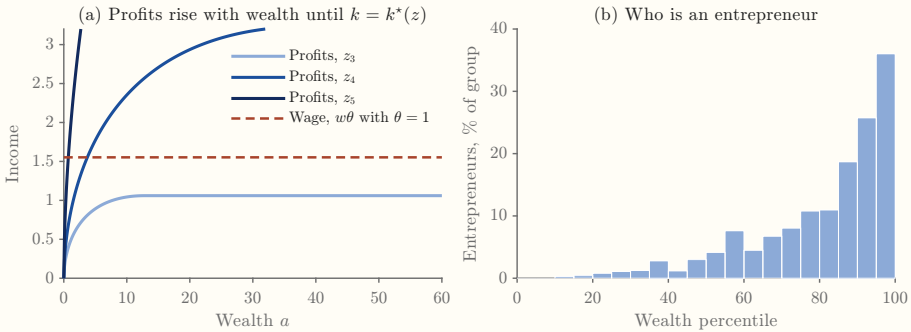
$$\pi(a,z) = \max_{k \le \phi a,\; l} \; z\left(k^{\alpha} l^{1-\alpha}\right)^{1-\nu} - w\,l - (r+\delta)\,k.$$

- Two features: **decreasing returns** cap the size of each firm; a **collateral constraint**, $k \le \phi a$, ties it to the owner's wealth.
- The budget constraint becomes $c + q\,a' = a + \max\{w\,\theta,\ \pi(a,z)\}$.
- A corporate sector, with the Aiyagari technology, hires the remaining capital and labor, so $w = w(r)$ as before.

Entrepreneurs: what changes in the algorithm

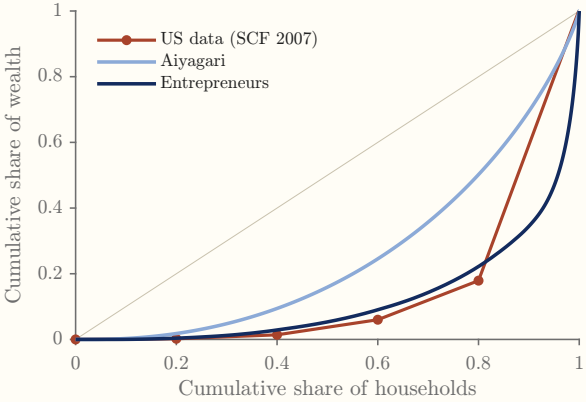
- Income $\max\{w\,\theta, \pi(a,z)\}$ has a **convex kink** where the household switches to entrepreneurship, so V need not be concave in a.
- Then the Euler equation is necessary but **not sufficient**: EGM may return a point that satisfies it but is not optimal, as warned in Part 2.
- Options:
  - VFI with Howard steps: it does not need concavity. We use it here, with $N_a = 300$ and 35 income–ability states.
  - An extension of EGM that compares the candidate points *(Iskhakov et al., 2017)*.
- The search over r is unchanged: bisection. The capital market now clears with $q \int a^\star \mathrm{d}\lambda = \int k \,\mathrm{d}\lambda + K_c$.
- A lesson: the right algorithm depends on the model, and robust methods are worth keeping in the toolbox.

Prediction: who becomes an entrepreneur



- Profits rise with wealth until the firm reaches its unconstrained size. Able but poor households must work, or run small firms, until they **save their way out** of the constraint.
- In a back-of-the-envelope calibration, entrepreneurs are 7% of households, hold 29% of wealth, run 35% of capital, and are 36% of the richest 5% (data: 14% of families, 40% of wealth).

Prediction: wealth inequality



- Entrepreneurs bend the Lorenz curve toward the data, mostly by emptying the middle.
- In this calibration, the wealth Gini is 0.48 without entrepreneurs, 0.75 with them and 0.78 in the data; the top 1% share is 4.7%, 22.9% and 33.5%.
- The interest rate falls, from 2.12% to -0.55%.
- What happens as collateral constraints loosen, and why the interest rate falls: Problem Set 2.

Summary

- In a stationary equilibrium with idiosyncratic risk, the distribution is constant, and equilibrium reduces to one price that clears one market.
- To compute it, search over r and, for each r, solve the household problem and find the distribution. Use Rouwenhorst for persistent income, EGM for the household, the histogram method for the distribution, and bisection for the price.
- Self-insurance lowers the interest rate (Huggett) and raises capital (Aiyagari), but barely moves the saving rate unless earnings risk is large and persistent.
- The benchmark predicts wealth more unequal than earnings, but less concentrated at the **top** than in the data, and too few households near the limit, hence a low average MPC.
- With entrepreneurs who face collateral constraints, the model predicts a top tail close to the data.
- Next lecture: transitional dynamics and aggregate risk.

References

Aiyagari, S. R. (1994). “Uninsured Idiosyncratic Risk and Aggregate Saving.” *Quarterly Journal of Economics* 109(3): 659–684.

Carroll, C. D. (2006). “The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems.” *Economics Letters* 91(3): 312–320.

Huggett, M. (1993). “The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies.” *Journal of Economic Dynamics and Control* 17(5–6): 953–969.

Iskhakov, F., T. H. Jørgensen, J. Rust and B. Schjerning (2017). “The Endogenous Grid Method for Discrete-Continuous Dynamic Choice Models with (or without) Taste Shocks.” *Quantitative Economics* 8(2): 317–365.

Judd, K. L. (1992). “Projection Methods for Solving Aggregate Growth Models.” *Journal of Economic Theory* 58(2): 410–452.

Kaplan, G. and G. L. Violante (2014). “A Model of the Consumption Response to Fiscal Stimulus Payments.” *Econometrica* 82(4): 1199–1239.

Kopecky, K. A. and R. M. H. Suen (2010). “Finite State Markov-Chain Approximations to Highly Persistent Processes.” *Review of Economic Dynamics* 13(3): 701–714.

References (continued)

Krueger, D., K. Mitman and F. Perri (2016). “Macroeconomics and Household Heterogeneity.” In *Handbook of Macroeconomics*, vol. 2A, 843–921. Elsevier.

Mehra, R. and E. C. Prescott (1985). “The Equity Premium: A Puzzle.” *Journal of Monetary Economics* 15(2): 145–161.

Quadrini, V. (2000). “Entrepreneurship, Saving, and Social Mobility.” *Review of Economic Dynamics* 3(1): 1–40.

Rouwenhorst, K. G. (1995). “Asset Pricing Implications of Equilibrium Business Cycle Models.” In T. F. Cooley (ed.), *Frontiers of Business Cycle Research*, 294–330. Princeton University Press.

Tauchen, G. (1986). “Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions.” *Economics Letters* 20(2): 177–181.

Uhlig, H. (1996). “A Law of Large Numbers for Large Economies.” *Economic Theory* 8(1): 41–50.

Zeldes, S. P. (1989). “Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence.” *Quarterly Journal of Economics* 104(2): 275–298.