Module 9 — Inequality and Distribution

Piketty, Markov Chains, and the Wealth Distribution

Dr Juan Zurita

From Smith to Simulation
The University of Edinburgh · School of Economics

Part I — The History

The Return of Inequality

- For most of the 20th century, mainstream economics had surprisingly little to say about inequality.
- The prevailing view, following **Simon Kuznets** (1955): inequality would take care of itself — rise during industrialisation, then fall.
- Starting around 1980, inequality began rising sharply in the English-speaking world.
- Top 1% income share in the US: fell from 24% (1928) to under 9% (1970s), then climbed back to over 20% by 2012.
- The comfortable Kuznets narrative broke down.

Piketty's *Capital in the Twenty-First Century* (2014)

"When the rate of return on capital exceeds the rate of growth of output and income . . . capitalism automatically generates arbitrary and unsustainable inequalities that radically undermine the meritocratic values on which democratic societies are based."

— Thomas Piketty, *Capital in the Twenty-First Century* (2014)

- A French economist publishes a 700-page book that becomes a global bestseller.
- Two centuries of tax data from over twenty countries.
- Central claim: the mid-20th-century decline in inequality was **the exception, not the rule**.

The Historical Arc: Three Eras

The Gilded Age (1870–1914) Extreme concentration of wealth — Rockefellers, Carnegies, Vanderbilts. The top 10% owned over 80% of total wealth in Europe and America.

The Great Compression (1930–1975) Two world wars, the Great Depression, progressive taxation, the rise of the welfare state dramatically *compressed* the wealth distribution. The anomaly Kuznets mistook for an iron law.

The Great Divergence (1980–present) Tax cuts, financial deregulation, globalisation, skill-biased technological change reversed the compression. Wealth concentration approached Gilded Age levels.

The Kuznets Curve

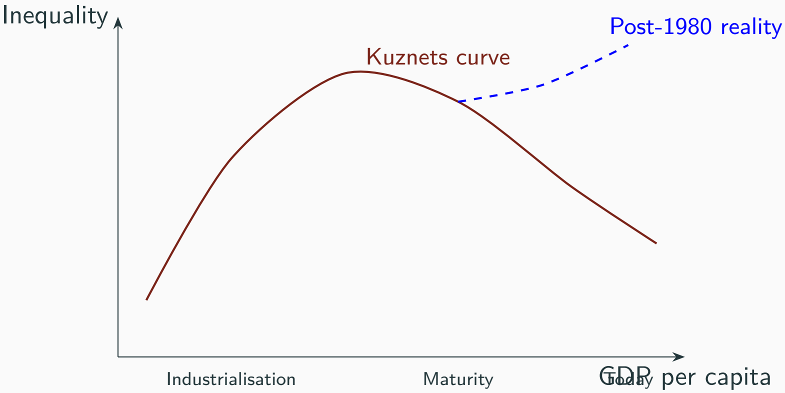

Kuznets (1955) predicted the inverted U. Piketty showed inequality can **rise again** — the natural tendency of capitalism when $r > g$.

Piketty's Central Formula: $r > g$

The Inequality

$$r > g$$

- r = rate of return on capital (dividends, interest, rents, capital gains)
- g = growth rate of the economy (GDP growth)

When $r > g$: inherited wealth grows **faster** than earned income. The rich get richer, not through effort, but through the *mathematics of compound returns*.

Historical Values

Era	r	g
1870–1914	5%	1.5%
1914–1950	1%	1.5%
1950–2012	5%	3.5%
2012–2100?	4%	1.5%

Aiyagari (1994): Inequality from Identical People

- Can we explain inequality *without* assuming people are fundamentally different?
- **S. Rao Aiyagari** (1994): yes.
- All agents are *identical* — same preferences, abilities, initial conditions.
- The only difference is luck: random income shocks modelled as a Markov chain.
- The wealth distribution that emerges is highly unequal and right-skewed.

Key Insight

Inequality is an emergent property of an economy with incomplete markets and idiosyncratic risk — not necessarily the result of differences in talent, effort, or inheritance.

Edinburgh Connection: Angus Deaton

- **Angus Deaton** (1945–2025), born in Edinburgh, educated at Fettes College.
- Nobel Prize in Economics, 2015.
- His book *The Great Escape* (2013): humanity has made extraordinary progress, but it has been **deeply uneven**.
- Developed household survey methods and consumption-based measures that let economists *see* inequality where GDP averages hid it.

Scottish Roots

The same empirical rigour and moral seriousness that characterised the Scottish Enlightenment — the insistence that economics must grapple with real data about real lives, not just elegant abstractions.

From History to Computation

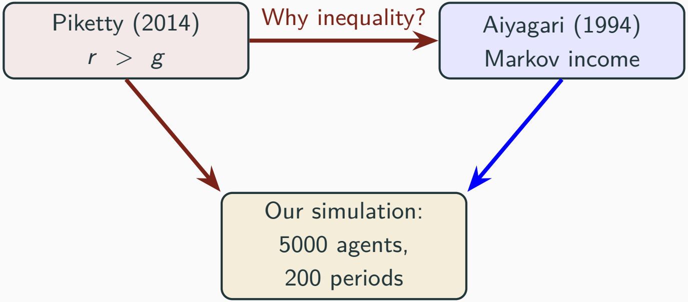

Lorenz curves, Gini coefficients, comparative statics

Part II — The Computation

Setting Up

```
import numpy as np
import matplotlib.pyplot as plt

np.random.seed(42)  # for reproducibility
```

- `numpy` — numerical arrays, random number generation, linear algebra
- `matplotlib` — plotting distributions and Lorenz curves
- No optimisation library needed this week — the model is **simulated**, not solved analytically

Income as a Markov Chain

Each agent's income follows a **Markov chain** — next period's income depends *only* on this period's income.

State	Label	Income y
0	Low	0.5
1	Medium	1.0
2	High	2.0

The **transition matrix** P:

$$P = \begin{pmatrix} 0.6 & 0.3 & 0.1 \\ 0.2 & 0.6 & 0.2 \\ 0.1 & 0.3 & 0.6 \end{pmatrix}$$

Entry P_{ij} = probability of moving from state i to state j. Each row sums to 1.

The Markov Chain: Transitions

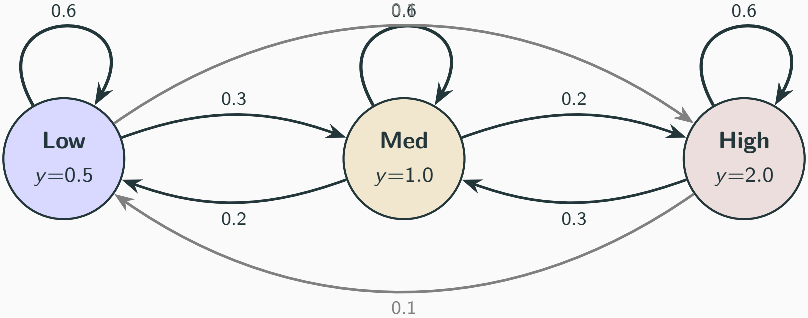

Moderate persistence: agents *tend* to stay in their current state but can move up or down. A low-income agent has a 60% chance of staying low, 30% of rising to medium, 10% of jumping to high.

Transition Matrix in Python

```python
# Income levels for three states
income_levels = np.array([0.5, 1.0, 2.0])

# Transition matrix: P[i,j] = Prob(state j | state i)
P = np.array([
    [0.6, 0.3, 0.1],  # from Low
    [0.2, 0.6, 0.2],  # from Medium
    [0.1, 0.3, 0.6],  # from High
])

print("Row sums:", P.sum(axis=1))  # [1. 1. 1.]
```

Key properties of any transition matrix:

- All entries ≥ 0 (probabilities)
- Each row sums to 1 (must go *somewhere*)
- Diagonal entries $=$ **persistence**

The Stationary Distribution

A Markov chain has a **stationary distribution** π such that $\pi P = \pi$.

- π tells us the long-run fraction of the population in each state.
- Found by computing the left eigenvector of P with eigenvalue 1.

For our transition matrix:

$$\pi = \left(\pi_{\text{Low}},\ \pi_{\text{Med}},\ \pi_{\text{High}}\right) \approx (0.235,\ 0.412,\ 0.353)$$

Long-run mean income: $\bar{y} = 0.235 \times 0.5 + 0.412 \times 1.0 + 0.353 \times 2.0 \approx 1.24$

Computing the Stationary Distribution

```python
# Solve pi @ P = pi  <=>  left eigenvector with eigenvalue 1
eigenvalues, eigenvectors = np.linalg.eig(P.T)

idx = np.argmin(np.abs(eigenvalues - 1.0))
pi = np.real(eigenvectors[:, idx])
pi = pi / pi.sum()  # normalise

for state, prob in zip(['Low', 'Medium', 'High'], pi):
    print(f"  {state}: {prob:.4f}")

mean_income = np.dot(pi, income_levels)
print(f"Long-run mean income: {mean_income:.4f}")
```

Simulating 5000 Agents

```
def simulate_markov(P, n_agents, n_periods, seed=42):
    rng = np.random.default_rng(seed)
    n_states = P.shape[0]
    states = np.zeros((n_agents, n_periods), dtype=int)
    states[:, 0] = rng.choice(n_states, size=n_agents, p=pi)
    for t in range(1, n_periods):
        u = rng.random(n_agents)
        for s in range(n_states):
            mask = (states[:, t-1] == s)
            cum_probs = np.cumsum(P[s])
            states[mask, t] = np.searchsorted(cum_probs, u[mask])
    return states

n_agents, n_periods = 5000, 200
state_paths = simulate_markov(P, n_agents, n_periods)
income_paths = income_levels[state_paths]
```

Each agent follows the same Markov chain — the only difference is luck.

From Income to Wealth: The Accumulation Equation

Agents accumulate wealth according to:

$$w_{t+1} = (1 + r) \cdot w_t + y_t - c_t$$

- w_t = wealth at time t
- r = interest rate (return on savings)
- y_t = income (from the Markov chain)
- c_t = consumption

Consumption rule: agents consume a fixed fraction α of resources:

$$c_t = \alpha \cdot \big((1 + r) \cdot w_t + y_t\big), \qquad \alpha \in (0, 1)$$

Borrowing constraint: $w_t \geq 0$ — agents cannot go into debt.

This friction is the key to the Aiyagari model: agents *cannot perfectly insure* against bad income shocks.

Wealth Simulation in Python

```python
def simulate_wealth(income_paths, r=0.02, alpha=0.9, w0=1.0):
    n_agents, n_periods = income_paths.shape
    wealth = np.zeros((n_agents, n_periods))
    wealth[:, 0] = w0
    for t in range(n_periods - 1):
        resources = (1 + r) * wealth[:, t] + income_paths[:, t]
        consumption = alpha * resources
        wealth[:, t+1] = np.maximum(resources - consumption, 0.0)
    return wealth

r = 0.02        # 2% interest rate
alpha = 0.9     # consume 90%, save 10%
wealth = simulate_wealth(income_paths, r=r, alpha=alpha, w0=1.0)
```

- All agents start with the same wealth $w_0 = 1$.
- After 200 periods: mean $>$ median $\Rightarrow$ right-skewed distribution.

The Stationary Wealth Distribution

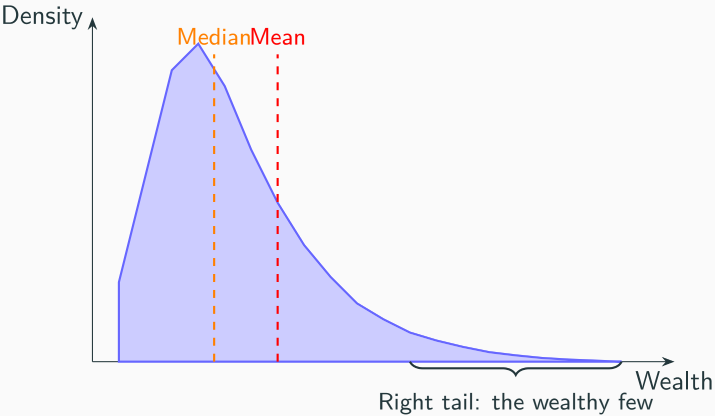

Mean > median: a few very wealthy agents pull the average up.

Identical agents, different luck $\Rightarrow$ unequal outcomes. This is Aiyagari's insight.

The Lorenz Curve

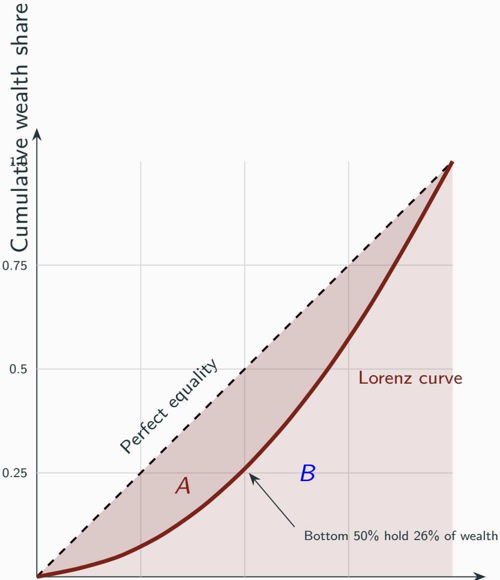

Computing the Lorenz Curve

```python
def lorenz_curve(wealth_array):
    sorted_wealth = np.sort(wealth_array)
    n = len(sorted_wealth)
    cumulative_wealth = np.cumsum(sorted_wealth)
    total_wealth = sorted_wealth.sum()
    # Prepend zero for the origin
    pop_share = np.concatenate(([0], np.arange(1, n+1) / n))
    wealth_share = np.concatenate(
        ([0], cumulative_wealth / total_wealth))
    return pop_share, wealth_share

pop_share, wealth_share = lorenz_curve(final_wealth)
plt.plot([0,1], [0,1], 'k--', label='Perfect equality')
plt.plot(pop_share, wealth_share, label='Lorenz curve')
plt.fill_between(pop_share, wealth_share, pop_share,
                 alpha=0.15)
```

The Gini Coefficient

Formula

$$G = 1 - 2\int_0^1 L(p)\,dp$$

- $G = 0$: everyone has the same wealth
- $G = 1$: one person has everything

Real-World Wealth Gini

Country	Gini
Sweden	≈ 0.50
Germany	≈ 0.67
UK	≈ 0.73
United States	≈ 0.85

Wealth Gini coefficients are much higher than income Gini coefficients — wealth accumulates over a lifetime and across generations.

Computing the Gini Coefficient

```python
def gini_coefficient(wealth_array):
    pop_share, wealth_share = lorenz_curve(wealth_array)
    # Area under Lorenz curve (trapezoidal rule)
    area_under_lorenz = np.trapz(wealth_share, pop_share)
    return 1 - 2 * area_under_lorenz

gini = gini_coefficient(final_wealth)
print(f"Simulated Gini: {gini:.4f}")
```

Wealth Shares (Typical Simulation Output)

- Top 1% own ∼5–8% of total wealth
- Top 10% own ∼25–35% of total wealth
- Bottom 50% own ∼20–30% of total wealth

Comparative Statics: What Drives Inequality?

Scenario	r	Income range	Effect on Gini
Baseline	0.02	[0.5, 1.0, 2.0]	—
Higher interest rate	0.05	[0.5, 1.0, 2.0]	$\uparrow$
Higher volatility	0.02	[0.2, 1.0, 3.0]	$\uparrow$

- **Higher** r: the rich earn more on their savings $\Rightarrow$ Piketty's $r > g$ mechanism.
- **Greater income volatility**: bigger shocks, harder to self-insure $\Rightarrow$ wider wealth dispersion.
- Both push the Lorenz curve **further from the line of equality**.

Running the Comparative Statics

```
# Scenario 2: Higher interest rate
wealth_high_r = simulate_wealth(income_paths, r=0.05,
                                alpha=0.9, w0=1.0)
gini_high_r = gini_coefficient(wealth_high_r[:, -1])

# Scenario 3: Higher income volatility
income_levels_volatile = np.array([0.2, 1.0, 3.0])
income_paths_volatile = income_levels_volatile[state_paths]
wealth_volatile = simulate_wealth(income_paths_volatile,
                                  r=0.02, alpha=0.9, w0=1.0)
gini_volatile = gini_coefficient(wealth_volatile[:, -1])
```

Plot all three Lorenz curves on one figure to visualise how parameter changes shift the distribution.

Lorenz Curves: Comparative Statics

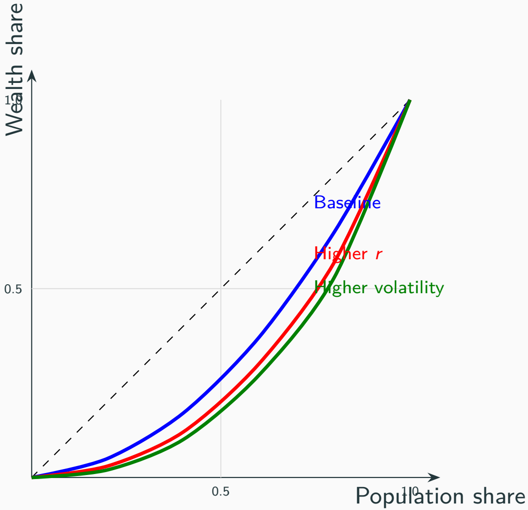

More bowed = more unequal. Both higher r and higher volatility **increase** the Gini coefficient.

How the Gini Evolves Over Time

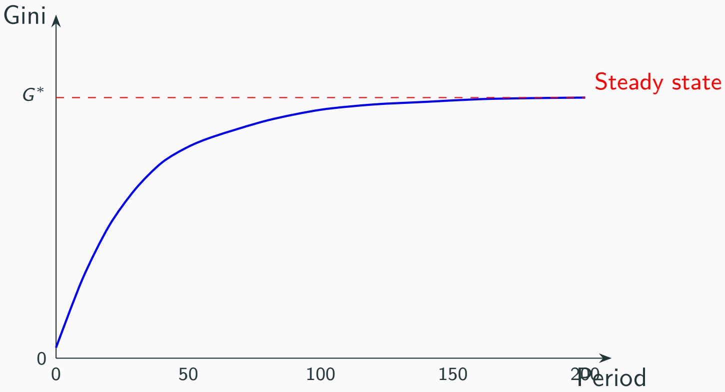

All agents start with $w_0 = 1$ (Gini ≈ 0). As idiosyncratic shocks accumulate, inequality rises and **converges to a steady state**.

Inequality is endogenous — emerging from luck alone.

Part III — Exercises & Discussion

Exercise 1: Income Mobility and Persistence

The transition matrix P determines how *persistent* income shocks are.

(a) Create a **high-persistence** transition matrix (diagonal $= 0.8$, off-diagonals split equally). Simulate 5000 agents for 200 periods. Compute the Gini.

(b) Create a **low-persistence** matrix (diagonal $= 0.4$). Simulate and compute the Gini.

(c) Plot the Lorenz curves for high-persistence, low-persistence, and baseline on a single figure. Which is most unequal? Why?

(d) Plot Gini vs. persistence for diagonal values $\{0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9\}$.

Intuition

Higher persistence $\Rightarrow$ long runs of bad (or good) luck $\Rightarrow$ larger wealth differences $\Rightarrow$ **more inequality**.

Exercise 2: Progressive Taxation

Implement a progressive tax-and-transfer system:

$$\text{tax}(y) = \begin{cases} 0.1 \cdot y & \text{if } y \leq 1.0 \\ 0.1 + 0.4 \cdot (y - 1.0) & \text{if } y > 1.0 \end{cases}$$

(a) Compute total tax revenue per period; redistribute equally as a lump-sum transfer. Simulate with post-tax income $y_t^{\text{net}} = y_t - \text{tax}(y_t) + \text{transfer}$.

(b) Compare Gini with and without taxation. Plot both Lorenz curves.

(c) Raise the top rate from 0.4 to 0.6. How much further does the Gini fall? Diminishing returns?

(d) Plot wealth distributions for no-tax, moderate-tax, and high-tax. Discuss the equality–efficiency trade-off.

Exercise 3: Piketty's $r > g$

(a) Add economic growth: $y_t = y_0 \cdot (1 + g)^t$. Simulate with $r = 0.04$, $g = 0.02$ ($r > g$). Track the Gini over 500 periods.

(b) Simulate with $r = 0.02$, $g = 0.04$ ($r < g$). Compare the two Gini trajectories on one plot.

(c) For the $r > g$ case, compute the top 10% wealth share over time. Does it stabilise or keep growing?

(d) Simulate Piketty's historical narrative:

- Periods 1–200: $r > g$ (Gilded Age)
- Periods 200–350: $r < g$ (Great Compression)
- Periods 350–500: $r > g$ (Great Divergence)

Plot the Gini over all 500 periods. Does it resemble the historical **U-shape** of inequality?

Quiz

Quiz: Conceptual Questions (1/2)

Q1. Piketty's $r > g$ means that when r exceeds g:

(a) The economy enters recession (b) Inherited wealth outpaces earned income

(c) Inflation accelerates (d) Government debt becomes unsustainable

Q2. A Gini coefficient of 0 means:

(a) Everyone has zero wealth (b) One person has all the wealth

(c) Wealth is perfectly equally distributed (d) The economy is in autarky

Q3. The Lorenz curve plots:

(a) Income over time (b) Cumulative wealth share vs. cumulative population share

(c) GDP growth vs. inflation (d) Supply against demand

Quiz: Conceptual Questions (2/2)

Q4. In the Aiyagari model, inequality arises because:

(a) Agents have different preferences
(b) Some agents are inherently smarter
(c) Identical agents face uninsurable idiosyncratic shocks
(d) The government redistributes from poor to rich

Q5. Piketty argued that the mid-20th-century decline in inequality was:

(a) The natural tendency of capitalism
(b) An anomaly caused by wars, depression, and progressive taxation
(c) The result of superior economic theory
(d) A statistical illusion

Quiz: Computational Questions (1/2)

Q6. If the Lorenz curve passes through (0.5, 0.2):

(a) The poorest 50% hold 20% of total wealth

(b) The richest 20% hold 50% (c) Average wealth is 0.2 (d) $G = 0.5$

Q7. A Gini of 0.85 (US wealth) indicates:

(a) Moderate inequality (b) Near-perfect equality

(c) Very high inequality — Lorenz curve bows far below equality

(d) 85% of people have equal wealth

Q8. The stationary distribution π of a Markov chain satisfies:

(a) $\pi = P$ (b) $\pi P = \pi$ (c) $\pi + P = 1$ (d) $\pi P = 0$

Quiz: Computational Questions (2/2)

Q9. Increasing the persistence of the income Markov chain will:

(a) Decrease inequality (more stable income)

(b) Increase inequality (long runs of bad/good luck)

(c) Have no effect (d) Make $G = 1$

Q10. The borrowing constraint $w_t \geq 0$ matters because:

(a) It prevents agents from saving

(b) It ensures all agents have the same wealth

(c) It prevents perfect insurance, generating inequality

(d) It makes the interest rate irrelevant

Key Takeaways

What We Learned This Week

1. **Piketty** (2014) showed that rising inequality is the historical norm; the mid-20th-century compression was the exception.
2. The formula $r > g$ captures the tendency for capital returns to outpace growth, concentrating wealth.
3. **Aiyagari** (1994) demonstrated that inequality can emerge from *identical* agents facing idiosyncratic shocks — inequality as an emergent property.
4. **Markov chains** model persistent income processes; Monte Carlo simulation generates wealth distributions.
5. The **Lorenz curve** and **Gini coefficient** are the standard tools for measuring inequality.
6. **Comparative statics**: higher r and greater income volatility both increase inequality.

Further Reading

- Piketty, T. *Capital in the Twenty-First Century* (2014), Introduction and Chs. 1, 7–12.
- Aiyagari, S.R. "Uninsured Idiosyncratic Risk and Aggregate Saving," *QJE* (1994).
- Deaton, A. *The Great Escape: Health, Wealth, and the Origins of Inequality* (2013).
- Piketty, T. and Saez, E. "Income Inequality in the United States, 1913–1998," *QJE* (2003).
- Kuznets, S. "Economic Growth and Income Inequality," *AER* (1955).

Next week: from the distribution of wealth *within* countries to the dynamics of growth *between* them —
convergence, divergence, and the puzzle of why some nations are rich and others poor.