Preface

These notes were assembled during the spring 2026 semester of the second-year PhD macroeconomics sequence at Penn State, taught by Maria-Jose Carreras-Valle (Part I) and Kai-Jie Wu (Part II). They aim to serve simultaneously as a compact reference for the technical machinery of modern macroeconomics—heterogeneous-agent equilibria, dynamic programming, business-cycle accounting, the empirics of consumption—and as a self-contained narrative of how the field's central questions evolve from one chapter to the next.

Audience and Prerequisites

The intended reader is a first- or second-year graduate student who has had a careful undergraduate or master's-level treatment of microeconomic theory (consumer choice, general equilibrium, basic dynamic programming) and the standard probability and real-analysis tools that come with that. No prior macroeconomics is strictly required, but the pace of *Part I* assumes familiarity with the Arrow–Debreu framework and the language of state-contingent claims.

Structure of the Book

The book is divided into two parts, reflecting the two-instructor structure of the course.

Part I: Heterogeneous Agents in Complete and Incomplete Markets (Chapters 1–3, by Maria-Jose Carreras-Valle) develops a unified framework for studying risk sharing across heterogeneous agents. Chapter 1 establishes the complete-markets benchmark—Arrow–Debreu trading, sequential trading, the recursive social planner—against which the rest of the book pushes. Chapter 2 introduces *exogenous* market incompleteness through Huggett, Aiyagari, and Krusell–Smith. Chapter 3 turns to *endogenous* incompleteness arising from participation frictions: one-sided lack of commitment, the Bulow–Rogoff model, and two-sided lack of commitment. The three chapters share a methodological signature: equilibria are characterized by the cross-sectional distribution of state variables, and the natural recursive formulation uses promised utility (or its analogue) as the state.

Part II: Growth, Business Cycles, and Quantitative Macroeconomics (Chapters 4–11, by Kai-Jie Wu) takes the dynamic-equilibrium machinery and applies it to canonical macroeconomic questions. Chapter 4 develops growth and development accounting as the empirical hook. Chapters 5–7 build the Solow and neoclassical growth models and confront them with cross-country convergence data. Chapter 8 extends to Real Business Cycles, and Chapter 9 inverts the RBC model to perform Business Cycle Accounting. Chapter 10 treats consumption and saving theory—the Permanent Income Hypothesis, Hall's Random Walk Hypothesis, and the empirical literature documenting excess sensitivity. Chapter 11 closes with the computation of the Aiyagari heterogeneous-agent model, which serves as the bridge into the modern HANK literature.

Pedagogical Conventions

Several typographic conventions recur throughout the text.

- **Definitions** appear in green-shaded boxes. **Theorems**, **Propositions**, **Lemmas**, **Corollaries**, and **Claims** appear in cyan-shaded boxes; their proofs follow inline (or in a dedicated grey-bordered block, when emphasized).
- **Remarks** come in two flavors. The shorter *inline* remarks (`\rmk`) flag a brief point in the surrounding narrative; the boxed *block* remarks (`\rmkb`) develop a substantial side topic, often spanning several paragraphs and including subsidiary figures or tables.
- **Algorithms** (e.g. Value Function Iteration, Aiyagari's outer loop) appear in violet-shaded boxes, listing the steps in order with implementation notes.
- **Examples** appear in their own environment with the worked solution clearly demarcated.
- **Facts** report empirical regularities in their own boxes, typically appearing in chapters that confront theory with data.

Each chapter opens with a brief *Notation in This Chapter* table listing chapter-specific symbols. The book-wide *Notation* section (immediately following this preface) collects symbols common to multiple chapters.

Reading Paths

Readers do not have to proceed linearly.

- *Heterogeneous-agent macro focus.* Read Part I in full, then Chapter 11 (Aiyagari computation). Chapter 10's PIH section provides useful background for the household problem in Aiyagari but is not strictly required.
- *Growth focus.* Read Chapters 4–7 as a self-contained block on growth theory and its cross-country evidence.
- *Business cycles focus.* Chapters 8–9 are the core; Chapter 10's RWH section complements the empirical discussion.
- *Computational focus.* Chapter 6 (Section on VFI), Chapter 8 (RBC numerical solution), and Chapter 11 (Aiyagari) form a sequence of progressively harder computational exercises.

Acknowledgments

These notes would not exist without Maria-Jose Carreras-Valle and Kai-Jie Wu, whose lectures form the underlying material. Any errors are mine—both as the typesetter and as the student.

Rui Zhou, Spring 2026

Notation

The following symbols recur throughout the notes. Where a chapter departs from a convention listed here, a chapter-specific note is provided in its opening section. A few high-level conventions:

- **Lowercase vs. uppercase letters.** Lowercase letters (e.g. c, k, y) denote per-worker or per-capita quantities. Uppercase letters (e.g. C, K, Y) denote aggregates. The convention is occasionally relaxed in specific chapters; when it matters, the chapter's notation note flags the exception.
- **Time subscripts.** t indexes the period; T is the terminal period in finite-horizon problems and the simulation length in numerical sections.
- **States and histories.** $s_t \in S$ is the period-t exogenous state; $s^t = (s_0, s_1, \ldots, s_t)$ is the history through date t.
- **Conditional expectation.** $\mathbb{E}_t[\cdot]$ denotes expectation conditional on the time-t information set.

Symbols used throughout the book.

Symbol	Meaning
Preferences and discounting	
$u(\cdot)$	Period utility function; $u' > 0$, $u'' < 0$, satisfying Inada conditions where needed.
β	Time discount factor; $\beta \in (0, 1)$.
σ	Coefficient of relative risk aversion under CRRA utility; the inverse $1/\sigma$ is the intertemporal elasticity of substitution.
γ	Coefficient of *absolute* risk aversion under CARA utility (Ch. 2 only).
$\mathbb{E}_t[\cdot]$	Expectation conditional on history s^t.
Stochastic environment	
s_t, s^t	Date-t state; history through t.
$\pi(s^t)$	Unconditional probability of history s^t; $\pi(s^\tau \| s^t)$ is conditional.
ε_t	Innovation / shock realization.
ρ	Persistence parameter of an AR(1) process; $\rho = \psi$ in Ch. 2's CARA example.
Endowment and production	
$y(s^t)$, Y_t	Stochastic endowment; aggregate output.

(continued on next page)

Symbol	Meaning
$F(K, L)$	Aggregate production function, typically constant returns to scale.
$f(k)$	Per-worker production function $f(k) = F(k, 1)$.
A, a_t	Total factor productivity (TFP); $a_t = \ln A_t$ for the log-linear AR(1) version.
α	Capital share in Cobb–Douglas production; output elasticity of capital.
δ	Depreciation rate of physical capital; $\delta \in (0, 1]$.
Quantities	
c, C	Consumption (per worker / aggregate).
k, K	Physical capital (per worker / aggregate).
L, l	Labor (aggregate / per worker). $L = 1$ in many setups.
I_t	Aggregate investment, $I_t = K_{t+1} - (1 - \delta)K_t$.
a, A	Asset / debt holdings (note: A is also used for TFP and natural debt limit; context disambiguates).
Prices and returns	
r	Real interest rate. Convention varies: in Ch. 1–3, 5–10, r is the net rate or rental rate of capital; in Ch. 11, $r = F_K(K, L)$ is the rental rate and the household's gross return is $1 + r - \delta$. Each chapter's notation note specifies the convention used.
R	Gross interest rate; typically $R = 1 + r$.
w	Real wage.
$q(s^t)$	Date-0 Arrow–Debreu price of a state-contingent claim (Ch. 1).
$Q(s'\|s)$	One-period-ahead pricing kernel in sequential trading (Ch. 1, 2).
Solution objects	
V	Value function.
$g(\cdot)$	Policy function.
Λ, λ	Cross-sectional distribution of agents (Ch. 2, 11).
Lagrangian and shadow prices	
$\mathcal{L}$	Lagrangian.
λ^i, μ^i	Pareto weight or Lagrange multiplier on a specific agent's budget; context distinguishes from the distribution λ.
$\theta(s^t)$	Multiplier on resource constraint (planner's problem, Ch. 1).
Empirical / decomposition objects	
$\mathrm{Var}, \mathrm{Cov}$	Cross-sectional variance and covariance.
g_x	Average growth rate of variable x over a sample period (Ch. 4).

A few overloaded symbols deserve attention. The Greek letter λ is used both for Pareto weights / Lagrange multipliers and for the cross-sectional distribution of agents—the role is always clear from context. The letter A is used for both the natural debt limit (Ch. 1) and TFP (Ch. 5 onward); these never appear together. The letter a is used for asset holdings throughout, and as log-TFP in Ch. 8; again no overlap.

Each chapter opens with a brief notation note flagging any chapter-specific symbols and confirming the local interpretation of r and a few other context-dependent objects.

Part I

Heterogeneous Agents in Complete and Incomplete Markets

Lectures by Maria-Jose Carreras-Valle

Part II

Growth, Business Cycles, and Quantitative Macroeconomics

Lectures by Kai-Jie Wu

Chapter 2

Computation of the Aiyagari Model

Remark (Notation in This Chapter).

Symbol	Meaning
z, z'	Idiosyncratic labor-productivity shock today and tomorrow
$\pi(z'\|z)$	Markov transition for the productivity shock
k, k'	Household's current / next-period capital holdings
B	Exogenous borrowing limit ($k' \geq -B$); $B = 0$ rules out borrowing
$V(z, k)$	Household's value function
$g_k(z, k), g_c(z, k)$	Optimal next-period capital and consumption
$\Lambda(z, k)$	Stationary cross-sectional distribution
$Q[(z, k), \cdot]$	Transition function on the state space
$X(r)$	Excess demand for capital, $K^{\text{firm}}(r) - \int g_k \, d\Lambda$
$K^{\text{firm}}(r), K^{\text{HH}}(r)$	Firm-side and household-side aggregate capital at candidate r
T (operator)	Bellman operator (VFI inner loop)
T (matrix)	Transition matrix on the discretized (z, k) state space

Convention on r: throughout this chapter r denotes the firm's rental rate of capital, $r = F_K(K, L)$. The household's gross return on saving is $1 + r - \delta$.

The previous chapters built the canonical *representative-agent* business cycle model and its accounting variant. By construction those models are silent on the **distribution** of wealth, consumption, or income across households—there is only one household. But a glance at any wealth-survey microdata file is enough to convince oneself that this distribution is far from a point mass: in the United States, the top 10% of households hold roughly 70% of net wealth, the bottom 50% hold under 2%. Any model that hopes to address questions about inequality, redistribution, social insurance, or precautionary saving must let agents differ in their wealth.

The **?** model is the workhorse for this purpose. It populates an otherwise standard neoclassical economy with a continuum of households who face uninsurable idiosyncratic income risk and a borrowing constraint, and it asks: what is the resulting cross-sectional distribution of wealth in the long run? Because this distribution is itself an endogenous object that feeds back into prices, solving the model is necessarily numerical. This chapter walks through the computation in detail, following the algorithm presented in lecture.

2.1 Why the Aiyagari Model?

The model captures four key features that, taken together, distinguish it from everything we have seen so far:

- **Ex-ante homogeneous households.** All households share the same preferences $u(c)$, the same discount factor β, and the same income process $\pi(z'|z)$. There is nothing about a household at birth that singles it out as "rich" or "poor."
- **Ex-post heterogeneous wealth.** Despite identical primitives, households end up with very different wealth holdings because they receive different histories of **idiosyncratic** income shocks $\{z_{i,t}\}$. A household that has drawn a long string of high-z realizations has accumulated savings; one that has drawn a long string of low-z realizations has run them down.
- **Incomplete financial markets.** There are no state-contingent claims that would allow households to insure against their idiosyncratic shock. The only available asset is a single risk-free bond / unit of capital. This is the friction: in a complete-markets economy (Chapter 1), the cross-sectional distribution of wealth would still be heterogeneous on the equilibrium path, but it would be Pareto-irrelevant—a planner could always reshuffle. Without state-contingent claims, individual histories matter for individual welfare.
- **Equilibrium concept: stationary wealth distribution.** The object of interest is not a sequence of allocations but the long-run *cross-sectional distribution* $\Lambda(z, k)$ that the economy converges to. The aggregate prices, r and w, are constants pinned down by the fixed point of this distribution.

Remark (The Trick: Continuum + iid $\Rightarrow$ No Aggregate Uncertainty).

> A continuum of households is not just notational convenience—it does real work. Because the shocks z_i are iid *across* households (though serially correlated within each household), a law of large numbers argument implies that aggregate quantities like the average labor supply $\int e^{z_i}\,di$ and the aggregate capital stock $\int k_i\,di$ are deterministic, even though every individual household is stochastic. This eliminates aggregate uncertainty and makes the equilibrium prices r and w *constants* in the steady state, which is what makes the problem tractable. With finitely many households, idiosyncratic shocks would never wash out and we would be back to a fully stochastic general-equilibrium problem (this is the Krusell–Smith setting we discuss at the end of the chapter).

Two Main Lessons

Fact 2.1: Aggregate vs. Idiosyncratic Risk: A New Perspective on the Equity Premium

? showed that calibrating a representative-agent model to match the historical equity premium ($\sim 6\%$ per year on stocks vs. $\sim 1\%$ on bonds) requires implausibly high coefficients of relative risk aversion ($\sigma \approx 30$). The puzzle is that aggregate consumption fluctuates very little, so a mildly risk-averse representative agent should not demand a large premium. Aiyagari-type models recast the premium as compensation not for smooth aggregate consumption risk but for highly volatile *idiosyncratic* risk that cannot be insured away. This shifts the explanation of the puzzle from preferences to market structure.

Fact 2.2: A Platform for Studying Redistribution and Insurance

Because the model has a non-degenerate wealth distribution as an endogenous object, it is the natural laboratory for analyzing policies that affect the distribution: progressive taxation, social insurance, unemployment insurance, public pensions, transfers. The welfare consequences of such policies cannot even be *stated* in a representative-agent framework. Almost every quantitative paper on redistribution since 1994 starts from some variant of Aiyagari.

The Wealth Distribution Tells the Story

To motivate the importance of the distribution, the figure below sketches the cross-sectional density of wealth that the model will deliver—right-skewed, with a mode near zero and a long right tail. This shape is qualitatively what we observe in U.S. data, although matching the very thick top tail is a known challenge for the basic model.

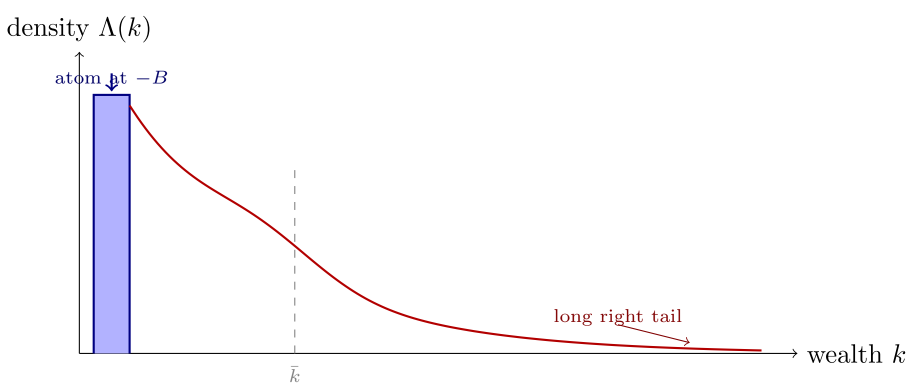

The atom at the borrowing constraint $-B$ is a generic feature of the model: a positive mass of households runs out of savings, hits the constraint, and stays there until a positive income shock pushes them off. We will see this atom emerge naturally from the algorithm.

2.2 The Environment

The setup is a stochastic, infinite-horizon, decentralized economy.

- **Time:** discrete, $t = 0, 1, 2, \dots$
- **Households:** a continuum indexed by $i \in I = [0, 1]$. All households share the same preferences
$$\mathbb{E}\left(\sum_{t=0}^{\infty} \beta^t\, u(c_{i,t}) \,\middle|\, z_{i,0}\right), \qquad u' > 0,\ u'' < 0.$$
- **Idiosyncratic income shock:** each household's labor productivity $z_{i,t}$ follows a Markov chain with transition $\pi(z' \mid z)$. Shocks are **iid across households** but **persistent within** each household. We will discretize z over a finite grid $\mathbb{Z} = \{z_1, \dots, z_{nz}\}$.
- **Asset:** the only available financial instrument is the economy's productive capital. Households save into it and earn a market return r. There are no state-contingent claims.
- **Borrowing constraint:** household assets must satisfy $k' \geq -B$, where $B \geq 0$ is an exogenous limit (e.g., $B = 0$ means no borrowing at all).
- **Firm:** a representative firm operates a constant-returns-to-scale technology $F(K, L)$ and rents capital and labor at competitive prices.

Remark (Why iid Shocks but Persistent Within a Household?).

> The two clauses are not contradictory. "iid across households" means: at any given date t, knowing household i's shock $z_{i,t}$ tells you nothing about household j's shock $z_{j,t}$. "Persistent within a household" means: knowing household i's shock today is informative about its shock tomorrow—the Markov chain has $\rho > 0$. Both features are necessary. Cross-sectional independence delivers the law of large numbers (no aggregate uncertainty). Within-household persistence delivers a non-trivial saving motive: if shocks were iid over time as well, the household would treat each draw as a transient surprise and save very little.

2.3 The Three Decentralized Problems

Because we are working with a decentralized economy (no planner), we write down the household, firm, and market-clearing conditions separately.

Household

Taking prices (r, w) as given, household i in state (z, k) chooses consumption and next-period assets:

$$\boxed{\begin{aligned} V(z,k) &= \max_{c,\,k'} \{u(c) + \beta\,\mathbb{E}(V(z',k') \mid z)\} \\ \text{s.t.}\quad & c + k' = (1 + r - \delta)\,k + w\,e^{z} \\ & k' \geq -B, \quad c \geq 0. \end{aligned}}$$

Note that capital depreciates at rate δ, so the gross return on a unit of capital is $1 + r - \delta$ (rental rate r minus depreciation, plus the principal back).

Remark ($1 + r$ vs. $1 + r - \delta$: A Common Notational Trap).

Different textbooks (and different photos of the lecture board) use different conventions. If r is interpreted as the *net interest rate on savings*, then the budget constraint reads $c + k' = (1 + r)k + we^z$ and the firm pays a rental rate $r + \delta$ to capital. If r is interpreted as the *rental rate on capital* (the firm's per-unit cost), then the budget constraint reads $c + k' = (1 + r - \delta)k + we^z$ since the household must absorb depreciation. We use the second convention throughout. Either is fine; what matters is consistency between the household's budget and the firm's FOC.

The two policy functions of interest are

$$g_k(z, k) \equiv \text{optimal } k', \qquad g_c(z, k) \equiv \text{optimal } c.$$

Remark (The Shape of g_k and Where the Borrowing-Constraint Atom Comes From).

The figure below shows the qualitative shape of $g_k(z, k)$ as a function of k, for two income levels $z_L < z_H$. Two features are essential for understanding the model.

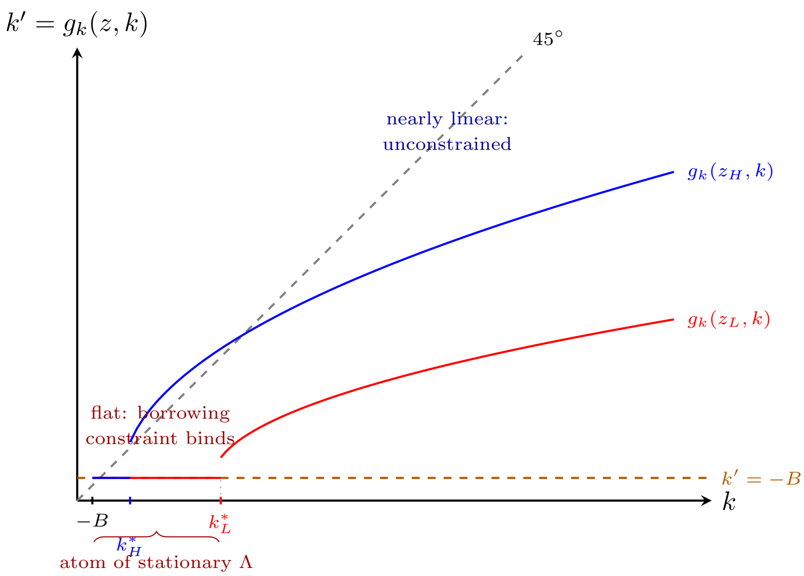

- **Flat segment at $k' = -B$.** For each income level z, there is a threshold k_z^* below which the household would optimally choose $k' < -B$ in the absence of the constraint. The constraint forces these households to set $k' = -B$ exactly. Households who land in the flat region therefore *accumulate at $k' = -B$*, which is the source of the atom in the stationary distribution $\Lambda(z, k)$ at the borrowing limit.
- **Income shifts the threshold.** A higher income $z_H > z_L$ provides more cushion against the next bad shock, so the household needs less assets to remain unconstrained. Hence $k_H^* < k_L^*$, and the high-income policy function leaves the flat region at a smaller

k.

- **Concavity above the threshold.** Once the constraint slackens, $g_k(z, k)$ is strictly concave in k and converges, for large k, to a near-linear schedule with slope close to $1/(1 + r)$. The richest agents are nearly behaving like permanent-income consumers: they save out of any windfall at a rate determined by the interest rate alone, regardless of their realized z. This is why aggregate capital is well-approximated by the rich-agent linear behavior even though the constrained-agent behavior is sharply non-linear.

The atom and the concavity together explain the right-skewed wealth distribution sketched at the start of this chapter: a positive mass piles up at the borrowing constraint (mostly low-z households), while persistently lucky households drift up the policy function and form the long right tail.

Firm

A representative firm rents aggregate capital K and labor L to maximize period profits:

$$\max_{K,\,L} F(K, L) - rK - wL.$$

Constant returns to scale plus competition imply zero profits and the standard FOCs

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

For a Cobb-Douglas $F(K, L) = K^\alpha L^{1-\alpha}$, this gives $r = \alpha(K/L)^{\alpha-1}$ and $w = (1-\alpha)(K/L)^\alpha$, which are functions of the capital-labor ratio alone. This will matter in Step 2 of the algorithm.

Market Clearing

Three markets must clear:

- **Labor:** aggregate labor supply equals labor demand, $L = \mathbb{E}(e^z)$ (the cross-sectional average of efficiency units, which is a constant by LLN).
- **Capital:** aggregate capital supply equals demand, $K = \mathbb{E}(k)$ (cross-sectional average wealth across households).
- **Goods:** $C + \delta K = F(K, L)$, where $C = \mathbb{E}(c)$ is aggregate consumption and δK is replacement investment (since aggregate K is constant in steady state).

2.4 Stationary Recursive Competitive Equilibrium

We can now define the equilibrium concept formally.

Definition 2.3: Stationary Recursive Competitive Equilibrium (SRCE)

A **stationary recursive competitive equilibrium** is a list

$$\{V(z,k),\ g_k(z,k),\ g_c(z,k);\ K,L;\ r,w;\ \Lambda(z,k)\}$$

consisting of a household value function and policy functions, firm choices of capital and labor, prices, and a probability measure Λ on $\mathbb{Z} \times \mathbb{K}$, such that:

(1) Given prices, V, g_k, g_c solve the household's problem;

(2) Given prices, (K, L) solve the firm's problem;

(3) The transition function induced by g_k and π leaves Λ invariant (stationarity, defined below);

(4) Markets clear:

$$L = \mathbb{E}(e^z)\,, \quad K = \int_{\mathbb{Z}\times\mathbb{K}} g_k(z,k)\, d\Lambda, \quad \int g_c(z,k)\, d\Lambda + \delta K = F(K,L).$$

2.5 The Distribution and Its Transition

The new object relative to the representative-agent model is the cross-sectional distribution Λ. Because both z and k are continuous (well, z is discrete after we discretize it; k is continuous in the model and on a fine grid in computation), we describe Λ as a probability measure on the product space $\mathbb{Z} \times \mathbb{K}$. For any (Borel) subset $\bar{Z} \subseteq \mathbb{Z}$ and $\bar{K} \subseteq \mathbb{K}$, $\Lambda(\bar{Z} \times \bar{K})$ is the mass of households whose state (z, k) lies in $\bar{Z} \times \bar{K}$.

The Transition Function Q

Given the household's policy g_k and the exogenous Markov chain π, the law of motion of (z, k) is mechanical:

$$z' \sim \pi(\cdot \mid z), \qquad k' = g_k(z,k).$$

The capital transition is deterministic conditional on the current state, while the income transition is stochastic. The **transition function** encodes this:

$$\boxed{Q\big[(z,k),\ \bar{Z}\times\bar{K}\big] \ \equiv\ \mathbf{1}\big\{g_k(z,k) \in \bar{K}\big\} \cdot \sum_{z' \in \bar{Z}} \pi(z' \mid z).}$$

The indicator captures that k' is pinned down deterministically by the policy; the sum captures the stochastic income transition. $Q[(z,k),\, \bar{Z} \times \bar{K}]$ is the probability that a household currently at (z, k) ends up in $\bar{Z} \times \bar{K}$ tomorrow.

Stationarity

The distribution Λ is stationary if it reproduces itself under Q:

$$\underbrace{\Lambda(\bar{Z} \times \bar{K})}_{\text{tomorrow's mass}} = \int_{\mathbb{Z}\times\mathbb{K}} \underbrace{Q\big[(z,k),\ \bar{Z} \times \bar{K}\big]}_{\text{prob. of landing in } \bar{Z}\times\bar{K}} \underbrace{d\Lambda}_{\text{today's mass}}$$

This is a fixed-point condition on Λ. Operationally, it says: tomorrow's mass in any region equals the integral, over all of today's (z, k), of the probability that a household at (z, k) moves into that region. It is exactly the cross-sectional analogue of the invariant distribution of a Markov chain.

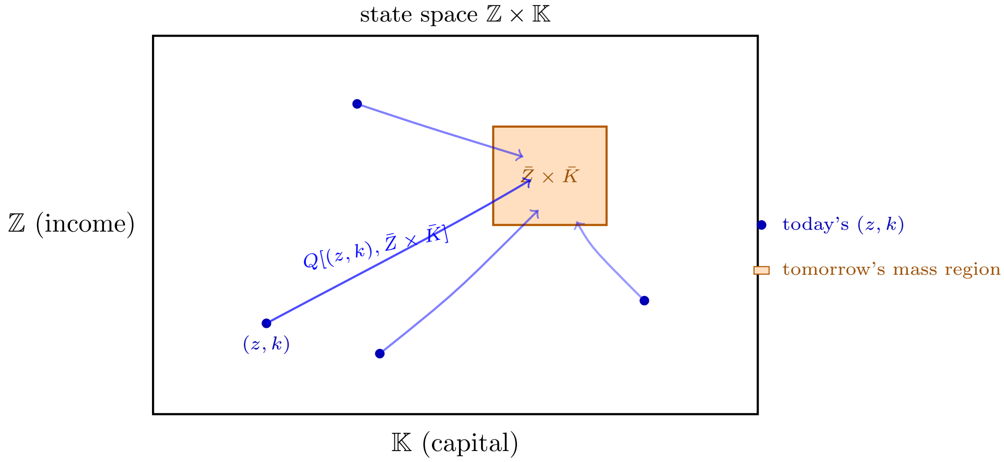

The diagram is the visual content of the stationarity equation: integrate the inflow probability Q over all today-states (blue dots), get tomorrow's mass in the orange box. Stationarity demands this equal $\Lambda(\bar{Z} \times \bar{K})$ itself.

2.6 The Computation Algorithm: Outer Loop on r

The equilibrium is jointly determined by $(V, g_k, g_c, \Lambda, r, w, K, L)$. The standard strategy is to nest the problem as follows:

- **Outermost loop**: search for the equilibrium interest rate r^*.
- Given a candidate r:
  - Recover w and the firm's K^{firm} from the firm FOCs.
  - Solve the household's Bellman equation by VFI to get g_k.
  - Iterate on the transition Q to find the stationary distribution Λ.
  - Compute the cross-sectional average wealth $K^{\text{HH}} = \int g_k \, d\Lambda$.
  - Form the excess demand $X(r) = K^{\text{firm}} - K^{\text{HH}}$.
- Adjust r until $|X(r)| < \varepsilon$.

Aiyagari Computation Algorithm (Outer Loop)

Step 1: **Guess** an initial r_0. A common choice is r_0 slightly below $1/\beta - 1$ (the representative-agent benchmark), since precautionary saving will push the equilibrium r^* below that value.

Step 2: **Compute w and K^{firm}.** Use the firm FOC $r_0 = F_K(K, L)$ together with the given $L = \mathbb{E}(e^z)$ to solve for the implied K/L ratio (and hence K^{firm}), then recover w from $w = F_L(K, L)$. With Cobb-Douglas:

$$\frac{K}{L} = \left(\frac{\alpha}{r_0}\right)^{\frac{1}{1-\alpha}}, \qquad w = (1-\alpha)\left(\frac{K}{L}\right)^{\alpha}.$$

Step 3: **Solve the household's problem.** Given (r_0, w), solve the Bellman equation by Value Function Iteration on a discretized (z, k) grid. The output is the policy function $g_k(z, k)$ (and $g_c(z, k)$, though the algorithm only needs g_k). This is the same VFI you implemented in PS3.

Step 4: **Solve for the stationary distribution** Λ given g_k. This is itself a fixed-point problem (an inner loop), described in Section 11.7.

Step 5: **Compute excess demand for capital:**

$$X(r_0) \;=\; \underbrace{K^{\text{firm}}(r_0)}_{\text{from Step 2}} \;-\; \underbrace{\int_{\mathbb{Z}\times\mathbb{K}} g_k(z, k)\, d\Lambda(z, k)}_{\text{from Steps 3 \& 4}}.$$

The integral is approximated on the discrete grid by $\sum_{i,j} g_k(z_i, k_j)\,\lambda(i, j)$, where λ is the discretized stationary distribution.

Step 6: **Convergence check:**

- If $|X(r_0)| < \varepsilon$, equilibrium is found. Output $r^* = r_0$.
- If $|X(r_0)| \geq \varepsilon$, update r_0 (e.g., by bisection if X has a known monotonicity; otherwise by a relaxation rule $r_0^{\text{new}} = r_0 - \kappa \cdot X(r_0)$ for small $\kappa > 0$) and return to Step 2.

The Equilibrium Picture: Capital Supply and Demand

The outer loop has a clean economic interpretation. Plot the candidate interest rate r on the vertical axis and the capital stock on the horizontal axis. There are two relations:

- **Capital demand** (downward sloping): the firm's $K^{\text{firm}}(r) = L \cdot (\alpha/r)^{1/(1-\alpha)}$, decreasing in r.
- **Capital supply** (upward sloping): the household sector's aggregate savings $K^{\text{HH}}(r) = \int g_k(z, k)\, d\Lambda(z, k; r)$. The supply schedule depends on r both through the policy function g_k (a higher return induces more saving) and through the stationary distribution Λ itself.

The equilibrium r^* is the value at which the two schedules cross. The excess-demand function $X(r)$ measures the horizontal gap:

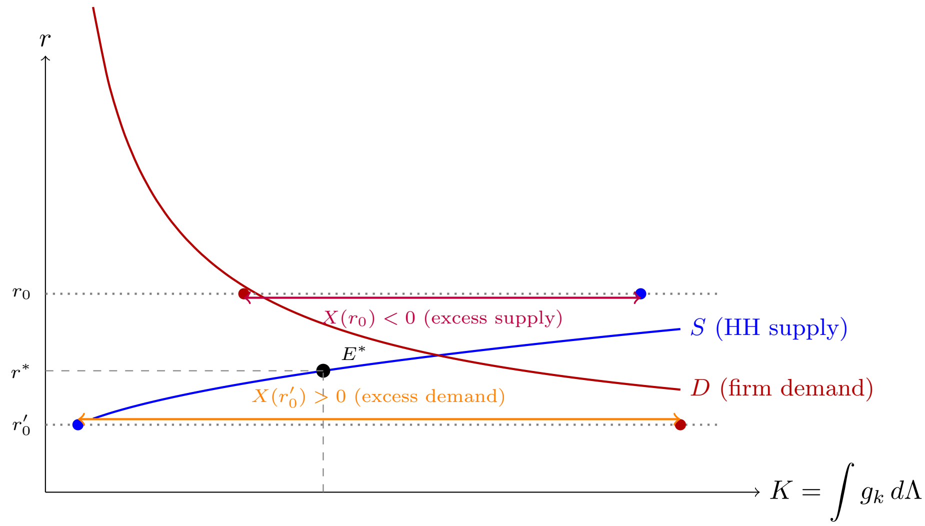

The picture clarifies the bisection logic: if at r_0 excess demand is negative (households want to save more than firms want to use), lower r_0; if it is positive, raise it.

Remark (Why $r^* < 1/\beta - 1$ in Aiyagari).

> In the deterministic representative-agent benchmark, the steady-state interest rate is pinned down by the Euler equation as $r^{\mathrm{RA}} = 1/\beta - 1$ (often written in net-of-depreciation form). In Aiyagari, $r^* < 1/\beta - 1$ *strictly*. The reason is precautionary saving: facing uninsurable income risk and a borrowing constraint, households accumulate more capital than the representative agent would, pushing K up and $r = F_K$ down. Quantitatively the gap is small (a few tens of basis points), but its existence is what gives the model its name and motivates the entire literature on incomplete markets. A subtler point: $\beta(1 + r^*) < 1$ is also *necessary* for the stationary distribution to exist—without it, individual wealth would drift off to infinity and there would be no fixed point of Λ.

2.7 Step 4 in Detail: Computing the Stationary Distribution

The inner step that deserves its own algorithm is Step 4: given a policy function g_k, find the Λ that satisfies

$$\Lambda(\bar{Z} \times \bar{K}) \;=\; \int_{\mathbb{Z}\times\mathbb{K}} Q\big[(z,k),\; \bar{Z}\times\bar{K}\big]\, d\Lambda.$$

On a continuous state space this is an integral fixed-point problem. We discretize.

Discretization

Replace $\mathbb{Z} \times \mathbb{K}$ by the finite grid

$$\{(z_i, k_j)\}_{i=1,\dots,n_Z;\; j=1,\dots,n_K}.$$

The distribution Λ becomes a probability matrix

$$\lambda \in \mathbb{R}^{n_Z \times n_K}, \qquad \lambda(i,j) = \Pr()\big(z = z_i,\ k = k_j\big), \qquad \sum_{i,j} \lambda(i,j) = 1.$$

The transition function Q becomes a 4-dimensional array indexed by today's (i, j) and tomorrow's (i', j').

The Update Rule

Given λ today, the mass at (i', j') tomorrow is

$$\boxed{\lambda_{\text{new}}(i', j') \;=\; \sum_{i=1}^{n_Z} \sum_{j=1}^{n_K} \lambda(i,j) \;\cdot\; \pi(i' \mid i) \;\cdot\; \mathbf{1}\{g_k(i,j) = j'\}\,.}$$

Reading the formula left to right: with mass $\lambda(i, j)$ at today's grid point (i, j), the income state moves to i' with probability $\pi(i' \mid i)$, and the capital state moves deterministically to whatever index j' corresponds to $g_k(i, j)$. Sum over all today-states to get tomorrow's mass at (i', j'). This is just Q applied as a linear operator on λ.

Stationary Distribution Algorithm (Step 4)

Step 4-1: Initial guess. Set $\lambda_0(i, j) = 1/(n_Z \cdot n_K)$ for all (i, j) (uniform). Any strictly positive distribution works; uniform is standard.

Step 4-2: Update. Compute

$$\lambda_1(i', j') = \sum_{i=1}^{n_Z} \sum_{j=1}^{n_K} \lambda_0(i,j)\, \pi(i' \mid i)\, \mathbf{1}\{g_k(i,j) = j'\}.$$

Step 4-3: Convergence check.

- If $\|\lambda_0 - \lambda_1\| < \varepsilon$ (e.g., sup-norm or ℓ^1 norm), stop. Output $\lambda^* = \lambda_1$.
- Otherwise, replace $\lambda_0 \leftarrow \lambda_1$ and return to Step 4-2.

Remark (The Iteration is Power Iteration on $Q^\top$).

Stack λ as a vector of length $n_Z \cdot n_K$. The update rule is then a linear map $\lambda_{\text{new}} = T\lambda$, where T is the matrix representation of the transition operator $Q^\top$ (the transpose because we are pushing the distribution forward, not pulling functions back). The stationary distribution is the dominant left-eigenvector of the Markov matrix—equivalently the right-eigenvector of T associated with eigenvalue 1. Power iteration finds it. For typical Aiyagari calibrations the convergence is geometric and reasonably fast (a few hundred iterations to machine precision), but degrades when the chain is nearly periodic or has near-unit-root subchains.

A Practical Caveat: Off-Grid Capital Choices

The clean formula $\mathbf{1}\{g_k(i,j) = j'\}$ assumes that the optimal k' lands exactly on a grid point. In practice it does not—$g_k(z_i, k_j)$ is some real number that falls between two grid values $k_{j'}$ and $k_{j'+1}$. Two standard fixes:

- **Nearest-neighbor.** Round g_k to the closest grid point and use the indicator as written. Easy, but introduces small biases that can compound over iterations.
- **Lottery (linear interpolation).** Split the mass between $k_{j'}$ and $k_{j'+1}$ in proportion to how close g_k is to each. If $g_k(i,j) = (1-\theta)k_{j'} + \theta k_{j'+1}$ with $\theta \in [0,1]$, then send a fraction $1-\theta$ of the mass to j' and θ to $j'+1$. This is the standard practice in the literature and is what most textbook implementations use under the hood.

Remark (Stationary Distribution as a Heatmap).

The output λ^* is most usefully visualized as a heatmap on the (z, k) grid. A typical pattern:

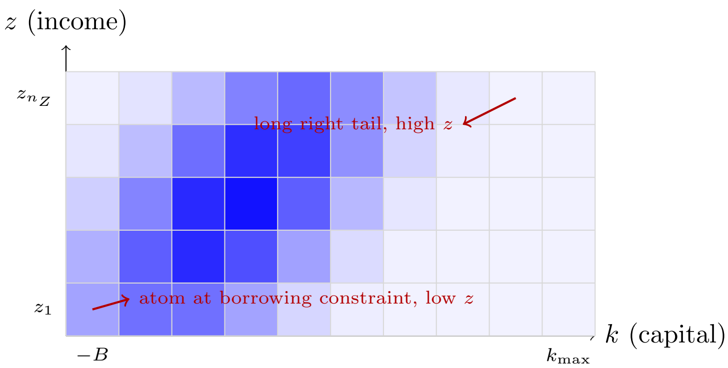

Mass concentrates along a positive ridge from low-z/low-k to high-z/high-k: persistently lucky households accumulate capital, persistently unlucky ones dissave to the borrowing limit. The marginal over k (summing across rows) is the wealth distribution figure from §1.

2.8 Visualizing the Algorithm

The full computation has *three* nested loops. It helps to keep them straight:

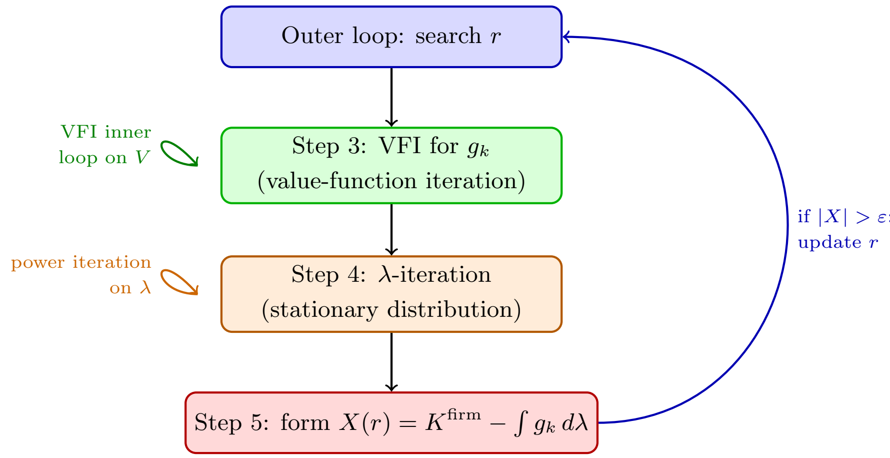

Three loops, three convergence criteria. The outer loop terminates when excess demand for capital is small. The VFI loop terminates when the value function stops moving. The λ-loop terminates when the distribution stops moving. Each loop sees the result of the inner loops as a fixed input.

Remark (Speedup Tricks Worth Knowing About).

The naive algorithm above works but can be slow. Standard speedups in the literature:

- **Howard improvement / policy iteration.** Once the policy g_k stops changing across VFI iterations, the value function update becomes a linear system that can be solved exactly rather than iterated. Cuts VFI time by a factor of 5–10x in typical cases.

- **Endogenous Grid Method (EGM, ? ?).** Avoids the expensive root-finding inside VFI by placing the grid on the post-decision asset k' rather than the pre-decision k, and inverting the Euler equation. Standard in modern HANK-style codes.

- **Direct eigenvector solve for λ.** Instead of iterating $\lambda_{\text{new}} = T\lambda$, just compute the dominant eigenvector of T directly via a sparse linear-algebra routine. Same answer, often 10x faster on moderately-sized grids.

2.9 Beyond the Basic Aiyagari Model

The basic model has shaped two decades of macro research, but several extensions are worth flagging.

- **? original quantitative findings.** With standard calibration, the equilibrium r^* is only a few tens of basis points below $1/\beta - 1$, and the wealth distribution generated by the basic model has a much thinner right tail than the U.S. data. Matching the top tail required innovations such as preference heterogeneity (Krusell–Smith 1998) or labor income with extreme tail risk (Castaneda–Díaz-Giménez–Ríos-Rull 2003).

- **Aggregate shocks: ?.** Adding an aggregate productivity shock a_t on top of idiosyncratic risk breaks the LLN trick—the distribution Λ now varies stochastically over time. Krusell and Smith showed that to a remarkable approximation, agents only need to track the aggregate capital stock K_t to forecast prices: "approximate aggregation." This made aggregate-shock heterogeneous-agent models computationally feasible.

- **HANK models.** Heterogeneous-Agent New Keynesian models combine an Aiyagari-style steady state with nominal rigidities (sticky prices, sticky wages) and aggregate shocks. The Aiyagari steady state is the starting point—all the new content concerns out-of-steady-state dynamics. Recent work (Auclert; Kaplan–Moll–Violante; McKay–Reis) shows that household heterogeneity reshapes the transmission of monetary and fiscal policy.

- **Continuous time methods. ?** reformulated Aiyagari in continuous time using Hamilton–Jacobi–Bellman equations and Kolmogorov forward equations, which often computes faster and gives crisper analytical results. Increasingly the modern standard.

Remark (The Big Picture).

> The Aiyagari model is a small step from the representative-agent neoclassical model—add idiosyncratic shocks, a borrowing constraint, and a continuum of agents—but it changes the questions one can ask. Distributional consequences of policy, the welfare cost of business cycles for the unlucky, the macroeconomic implications of inequality: all became quantitatively tractable. The computational machinery developed for Aiyagari (VFI, distribution iteration, the outer-loop price search, EGM) is the same machinery that powers HANK models and, more broadly, almost all of modern quantitative macro. Mastering the algorithm in this chapter is therefore not just a problem-set exercise but the entry ticket to a large fraction of the contemporary research frontier.

Remark (Chapter Summary).

> - **The three-loop computation.** Outer loop searches for r^* that clears the capital market; the middle loop is VFI for the household policy $g_k(z, k)$; the innermost loop iterates the wealth distribution $\Lambda(z, k)$ to its stationary fixed point.
>
> - **Where the wealth distribution shape comes from.** The borrowing constraint produces an atom at $k = -B$ (constrained households who would optimally borrow more). The right tail is generated by persistently lucky households drifting along the (concave but eventually near-linear) policy function g_k.
>
> - **Equilibrium $r^* < 1/\beta - 1$ strictly.** Precautionary saving raises aggregate household savings above the representative-agent benchmark, which depresses the equilibrium interest rate.
>
> - **Standard speedups matter.** Howard improvement, the Endogenous Grid Method (Carroll, 2006), and direct sparse-eigenvector solves for Λ each cut runtime by a substantial factor on realistic grids.
>
> - **Bridge to the modern frontier.** The Aiyagari machinery generalizes to Krusell–Smith (aggregate shocks), continuous-time HJB/KFE formulations (Achdou et al.),

and HANK models with nominal rigidities. Mastering this chapter unlocks much of contemporary quantitative macro.

Part III

Problem Sets and Solutions

Subject Index

Author Index