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; $s^t = (s_0, s_1, \dots, s_t)$ 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; $\pi(s^\tau \vert s^t)$ 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

Chapter 2

Exogenously Incomplete Markets

Remark (Notation in This Chapter).

Symbol	Meaning
a, a'	Today's and tomorrow's holdings of the single risk-free bond
$R = 1 + r$	Gross interest rate; partial-equilibrium parameter, GE-determined later
ϕ	Ad-hoc borrowing limit; potentially tighter than the natural debt limit
$x = y(s) + Ra$	"Cash-in-hand" (the i.i.d. case state-space reduction)
$\hat{a}, z$	Change-of-variables for $\phi < 0$: $\hat{a} \equiv a - \phi$, $z \equiv x - \phi$
ψ	Persistence of the AR(1) endowment process (CARA example)
σ^2	Variance of the AR(1) innovation ε
$\bar{y}$	Long-run mean of endowment
$\lambda(a, s)$	stationary distribution of (a, s) across agents
$g(a, s)$	Savings policy function $a' = g(a, s)$
z_t	Aggregate productivity shock (Krusell–Smith only)
$H(\lambda, z, z')$	Aggregate law of motion for the wealth distribution (Krusell–Smith)
$m = (m_1, m_2, \ldots)$	Moments parameterizing the distribution (Krusell–Smith approximation)

Unlike the complete market case, in the incomplete market case, agents cannot fully insure against idiosyncratic risk. This is because there are not enough assets to span all possible states of the world. Specifically, we assume agents can buy or sell only one kind of security: risk-free bonds.

2.1 Two-Period Example

Assume:

- Agents can buy or sell a risk-free bond.
- Partial equilibrium: The return of the bond, R ($R \geq 1$) is fixed and held constant.

- $t = 0$: y_0 is fixed and known.
- $t = 1$: $y(s)$ is stochastic with $y_1 < y_2 < \cdots < y_N$.

The problem of the agent is given by

$$\begin{aligned}\max_{c_0,\{c(s)\}_s,a} \quad & u(c_0) + \beta \sum_s \pi(s) u(c(s)) \\ \text{s.t.} \quad & c_0 + a \le y_0 \\ & c(s) \le y(s) + Ra, \quad \forall s \\ & a \ge -\frac{\min y(s)}{R} = -\frac{y_1}{R}.\end{aligned}$$

This borrowing constraint differs from the complete-markets natural debt limit: it requires the agent to be able to repay in *every* possible future state from a single risk-free bond, rather than in expectation across the menu of state-contingent claims.

We can rewrite the problem as

$$\begin{aligned}\max_a \; & u(y_0 - a) + \beta \sum_s \pi(s) u(y(s) + Ra) \\ \text{s.t.} \; & a \ge -\frac{y_1}{R}.\end{aligned}$$

Let λ be the Lagrange multiplier for the borrowing constraint. The FOC is given by

$$-u'(y_0 - a) + \beta R \sum_s \pi(s) u'(y(s) + Ra) + \lambda = 0,$$

Note that here $\lambda = 0$. If $\lambda > 0$, then the borrowing constraint is binding, and the agent will have zero consumption in the worst state. This is suboptimal by the Inada condition. Hence, $\lambda = 0$, and the FOC can be rewritten as

$$u'(y_0 - a) = \beta R \sum_s \pi(s) u'(y(s) + Ra).$$

Here we also consider the corresponding complete market case:[1]

$$\begin{aligned}\max_{c_0,\{c(s)\}_s,\{a(s)\}_s} \quad & u(c_0) + \beta \sum_s \pi(s) u(c(s)) \\ \text{s.t.} \quad & c_0 + \sum_s Q(s) a(s) \le y_0 \\ & c(s) \le y(s) + a(s), \quad \forall s.\end{aligned}$$

---

[1] In reality, the complete market also has a borrowing constraint, but it is "hidden" by the market structure and the agent's preferences. The logic is basically the same as previously argued in the incomplete market case. In complete market, the agent trades state-contingent claims $a(s)$. The obligation to repay is strictly tied to the specific state s that materializes. Therefore, the natural debt limit is state-specific: $c(s) = y(s) + a(s) \ge 0 \implies a(s) \ge -y(s)$. Standard macroeconomic utility functions satisfy the Inada condition ($\lim_{c\to 0} u'(c) = \infty$). A rational agent will never choose a consumption allocation where $c(s) = 0$ in any state s, because the marginal utility of that first unit of consumption is infinite. Consequently, the state-by-state natural debt limit $a(s) \ge -y(s)$ will *never bind* (we are mathematically guaranteed an interior solution). Because it is never binding, it is standard practice to omit it from the formal setup and treat the choice of $a(s)$ as unconstrained.

And this can be rewritten as

$$\max_{a(s)} u\left(y_0 - \sum_s Q(s)a(s)\right) + \beta \sum_s \pi(s)u(y(s) + a(s)).$$

Define

$$\hat{a} = \sum_s Q(s)a(s),$$
$$\beta = 1/R.$$

The FOC to the complete market case is given by

$$Q(s)u'(y_0 - \sum_s Q(s)a(s)) = \beta u'(y(s) + a(s)) \implies Q(s) = \beta\pi(s)\frac{u'(c(s))}{u'(c_0)}.$$

In equilibrium, in complete markets, $c(s) = c_0$ for all s, which implies that $Q(s) = \beta\pi(s)$ for all s.[2] Hence,

$$\hat{a} = \sum_s Q(s)a(s) = \beta \sum_s \pi(s)a(s) = \frac{\sum_s \pi(s)a(s)}{R} \implies R\hat{a} = \sum_s \pi(s)a(s).$$

So we can rewrite the complete market problem as[3]

$$\max_{\hat{a}} u(y_0 - \hat{a}) + \beta \sum_s \pi(s)u(y(s) + R\hat{a}).$$

The FOC then gives

$$u'(y_0 - \hat{a}) = \beta R u'\left(\sum_s \pi(s)y(s) + R\hat{a}\right) = \beta R u'\left(\mathbb{E}[y(s)] + R\hat{a}\right).$$

Recall that the FOC for the incomplete market case gives

$$u'(y_0 - a) = \beta R \sum_s \pi(s)u'(y(s) + Ra) = \beta R \mathbb{E}[u'(y(s) + Ra)].$$

If we further assume that $u' > 0$, $u'' < 0$ and $u''' > 0$, we can formally show that the optimal savings in the incomplete market is strictly greater than in the complete market:

$$a > \hat{a}.$$

The condition $u''' > 0$ implies that the marginal utility function $u'(\cdot)$ is strictly convex. By Jensen's Inequality, for any non-degenerate random variable $y(s)$, the expected marginal

---

[2] The result $c(s) = c_0$ stems from two distinct smoothing mechanisms. First, *cross-state smoothing*: complete markets allow the agent to fully insure against idiosyncratic risk, so consumption is constant across all states tomorrow ($c(s) = c_1$ for all s). Second, *intertemporal smoothing*: the assumption $\beta = 1/R$ (i.e., $\beta R = 1$) implies that the agent's impatience exactly offsets the market interest rate, so there is no incentive to tilt consumption across periods, yielding $c_0 = c_1$. Combining the two forces gives $c(s) = c_0$ for all s, which allows the marginal utilities to cancel out, leaving the actuarially fair state prices $Q(s) = \beta\pi(s)$.

[3] This is true since $c(s) = y(s) + a(s) := c_0$ is a constant across s, $\sum_s \pi(s)u(y(s) + a(s)) = \sum_s \pi(s)u(c(s)) = u(c_0) = u\left(\sum_s \pi(s)c(s)\right) = u(\sum_s \pi(s)y(s) + \sum_s \pi(s)a(s)) = u(\sum_s \pi(s)y(s) + R\hat{a})$

utility is strictly greater than the marginal utility of the expected consumption:

$$\mathbb{E}[u'(y(s) + R\hat{a})] > u'(\mathbb{E}[y(s)] + R\hat{a}).$$

If the agent saved the same amount $\hat{a}$ in both economies, the right-hand side of the incomplete-market Euler equation would be strictly larger. To restore optimality, the agent must increase a, which lowers tomorrow's expected marginal utility and raises today's marginal utility $u'(y_0 - a)$ until equality holds. Thus, $a > \hat{a}$.

Remark.

> - **Prudence:** An agent with $u''' > 0$ is said to be *prudent.* In complete markets, the agent insures away all idiosyncratic risk, so they only save for intertemporal smoothing. In incomplete markets, the agent is fully exposed to the variance of $y(s)$. Because they are prudent, this uninsurable uncertainty creates a *precautionary saving motive.* They build a "buffer stock" of wealth today to self-insure against the possibility of realizing a terrible shock tomorrow.
>
> - **Welfare Implication:** It is crucial to note that saving more does *not* mean the agent is better off. In fact, by the First Welfare Theorem, the complete market allocation is Pareto optimal. In the incomplete market, the agent suffers a welfare loss due to uninsurable risk and is forced to painfully sacrifice current consumption (c_0 drops) merely to build a defensive buffer against future volatility.

2.2 Huggett Model

The two-period example above made the central economic point—that incomplete markets generate precautionary saving relative to the complete-markets benchmark. We now extend the framework to an *infinite-horizon* setting and let the resulting cross-sectional distribution of wealth become an equilibrium object in its own right. The **?** model is the canonical formulation: a continuum of ex-ante identical households, hit by uninsurable idiosyncratic income shocks, with a single risk-free bond in zero net supply.

We develop the model in three steps. First, we treat the household's problem as a partial-equilibrium recursive program, taking the interest rate R as given. Second, we close the model in general equilibrium by requiring that the bond market clear at the equilibrium rate. Third, we describe how to compute the stationary cross-sectional distribution of wealth that the model produces. Each step will reappear—almost verbatim—in the Aiyagari and Krusell–Smith extensions later in this chapter, and again in the Aiyagari computation chapter at the end of the notes.

2.2.1 Partial Equilibrium

Assume:

- Interest rate: $R = 1 + r$ is fixed and held constant.

- Preferences:

$$\sum_t \sum_{s^t} \beta^t \pi(s^t) u(c(s^t)).$$

- Endowment:
  - $y(s_t) \in \{y_1, y_2, \cdots, y_N\}$ such that $y_1 < y_2 < \cdots < y_N$.
  - Endowment process is i.i.d.: $\pi_1, \pi_2, \cdots, \pi_N$.

Before formally writing down the Bellman equation, we must identify the state variables—the minimal set of information the agent needs to make an optimal decision today. In this incomplete market environment, the agent's situation is fully summarized by two variables, (a, s):

- **Endogenous State (a):** The agent's current asset holding. It acts as a summary statistic for their entire history of past shocks and consumption-saving decisions, determining the financial buffer they bring into today.

- **Exogenous State (s):** The realization of today's idiosyncratic shock. Because the endowment process is assumed to be i.i.d., today's state s provides absolutely no information about tomorrow's state s'. Its *only* role is to determine today's flow income $y(s)$.

Together, (a, s) fully pin down the agent's current purchasing power. The recursive problem is then given by

$$\begin{aligned} v(a, s) = \max \ & u(c(s)) + \beta \sum_{s'} \pi(s') v(a', s') \\ \text{s.t. } & c(s) + a' \leq y(s) + Ra, \quad \forall s \\ & a' \geq -\phi. \end{aligned}$$

Here ϕ is the *ad-hoc* borrowing constraint. Different from the natural borrowing limit, it might be the case that $\phi \leq \frac{\min\{y_1, \cdots, y_N\}}{R} = \frac{y_1}{R}$ (i.e., the ad-hoc borrowing constraint is more stringent than the natural borrowing limit). In this case, the ad-hoc borrowing constraint might be binding, and thus we cannot simply omit it from the setup.

In this problem, we can use "case-in-hands" to reduce dimension of the state space (since the states are i.i.d.):

$$\begin{aligned} & c + a' = y_s + Ra \\ \Longrightarrow\ & c = y_s + Ra - a' \\ \Longrightarrow\ & c = x - a' \quad \text{ where } x = y_s + Ra. \end{aligned}$$

Here $x = y_s + Ra$ captures the endowment, state s and asset a.

Remark (Why doesn't x lose any crucial information?).

> - **Economic Intuition of x:** The variable $x = y_s + Ra$ represents the total liquid wealth the agent has available to spend or save at the very beginning of the period. Once yesterday's savings (Ra) and today's income (y_s) are pooled together, a dollar from savings is indistinguishable from a dollar from today's labor.
>
> - **Why we can drop s here:** A state variable must tell us either about current resources or about future probabilities. Once s has realized and delivered today's income y_s (now absorbed into x), it has done its first job. Because the shock is i.i.d., today's s

carries no predictive content for tomorrow's shock s', so it is no longer needed to form expectations about the future. The state space collapses from (a, s) to just x without any loss of information.

- *Note:* If s were persistent (e.g., a Markov chain or AR(1) process), we would still combine resources into x, but we would have to keep s as a separate state variable to forecast tomorrow, yielding $v(x, s)$. The AR(1)-CARA example below is exactly such a case.

We can rewrite the problem as

$$\begin{aligned} v(x) &= \max_{a'} \; u(x - a') + \beta \sum_{s'} \pi(s') v(x') \\ \text{s.t. } & x \geq a' \\ & a' \geq -\phi. \end{aligned}$$

The FOC is given by

$$-u'(x - a') + \beta R \sum_{s'} \pi(s') v'(y(s') + Ra') + \lambda = 0,$$

where λ is the Lagrange multiplier for the borrowing constraint. Note that here we cannot argue that $\lambda = 0$ since the borrowing constraint is ad-hoc and might be different from the natural borrowing limit, and thus may be binding.

Since $\lambda \geq 0$, the FOC implies that

$$u'(x - a') \geq \beta R \sum_{s'} \pi(s') v'(y(s') + Ra') = \beta R \mathbb{E}[v'(x')].$$

By the envelope condition, we have $v'(x) = u'(x - a')$. Hence, we can rewrite the FOC as

$$v'(x) \geq \beta R \mathbb{E}[v'(x')].$$

We pause to establish a key result: assets diverge to infinity whenever $\beta R \geq 1$.

Lemma 2.1: Assets Diverge to Infinity When $\beta R \geq 1$

If $\beta R \geq 1$, then

$$\lim_{t \to \infty} x_t = \infty.$$

Proof for Lemma

The FOC is the same as the partial equilibrium case:

$$v'(x) \geq \beta R \mathbb{E}[v'(x')] \geq \mathbb{E}[v'(x')].$$

This implies that $v'(x)$ is a non-negative supermartingale. By Doob's Convergence Theorem, $v'(x)$ converges almost surely to a limit v'_∞ as $t \to \infty$.[a] Thus, for a sequence of endowments $\{y_t(s)\}$, the sequence $\{v'(x_t)\}$ converges to a particular number,

$$\lim_{t\to\infty} v'(x_t) = \tilde{v} \geq 0.$$

Claim

$$\lim_{t\to\infty} v'(x_t) = 0.$$

Proof for Claim.

> Suppose in contradiction that $\lim_{t\to\infty} v'(x_t) = \tilde{v} > 0$. This means x_t must converge to a finite positive value $(v')^{-1}(\tilde{v})$. But by construction $x_t = y(s) + Ra_t$, and $y(s)$ is i.i.d. Hence, x_t cannot converge to a finite positive value, which leads to a contradiction. Therefore, $\lim_{t\to\infty} v'(x_t) = 0$.

$\lim_{t\to\infty} v'(x_t) = 0 \implies \lim_{t\to\infty} x_t = \infty$. This means that the agent's assets diverge to infinity. The intuition is that the agent is patient enough so that they always have enough incentives to save for the future.

---

[a] A *supermartingale* is a sequence of random variables $\{X_t\}$ such that $\mathbb{E}[X_{t+1}|X_t, X_{t-1}, \ldots] \leq X_t$. Intuitively, given all past information, the expected future value does not exceed the current value. A key property of supermartingales is that if a supermartingale is also bounded below, it must converge. *Doob's Convergence Theorem* (also called Doob's Martingale Convergence Theorem) states that any non-negative supermartingale must converge almost surely to a finite limit. In our context, since $v'(x) \geq \mathbb{E}[v'(x')] \geq 0$, the sequence of marginal utilities forms a non-negative supermartingale, and thus must converge to some limiting value v'_∞ along almost all paths of endowment realizations.

2.2.2 CARA Example

Assume:

- $u(c) = -\frac{1}{\gamma}\exp(-\gamma c)$, where $\gamma > 0$ is the coefficient of absolute risk aversion.
  - $v(\cdot)$ then inherits properties of $u(\cdot)$: increasing, strictly concave and differentiable.
- Natural borrowing limit: $\phi = \frac{y_1}{R} \implies$ no need to keep track of the borrowing constraint.
- Endowment process: AR(1)[4][5]

$$y(s') = \psi y(s) + (1-\psi)\bar{y} + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \sigma^2), \quad |\psi| \in (0,1).$$

- Start by assuming $R > 1$.

---

[4] An *AR(1) process* (autoregressive process of order 1) is a stochastic process where the current value depends on the previous value plus a random shock. The general form is $y_t = \rho y_{t-1} + (1-\rho)\mu + \epsilon_t$, where $|\rho| < 1$ ensures stationarity, μ is the long-run mean, and ϵ_t is white noise. The parameter ρ (here denoted ψ) measures the persistence of the process: if ρ is close to 1, the process exhibits high persistence (shocks have long-lasting effects); if ρ is close to 0, the process is closer to white noise. AR(1) processes are commonly used in macroeconomics to model persistent endowment or income shocks.

[5] Note that here the endowment process is not i.i.d. as previously assumed. This means that today's state s does provide information about tomorrow's state s'. Therefore, we cannot drop s as a state variable, and the value function must be written as $v(x, s)$ instead of just $v(x)$.

The recursive problem is given by

$$\begin{aligned} v(x,s) &= \max_{a'} \ -\frac{1}{\gamma}\exp(-\gamma(x-a')) + \beta\sum_{s'}\pi(s'|s)v(y(s')+Ra',s') \\ &= \max_{a'} \ u(x-a') + \beta\mathbb{E}[v(y(s')+Ra',s')|s]. \end{aligned}$$

Note that here we do not keep track of the borrowing constraint since we have assumed that ϕ is the natural borrowing limit.

We use the **guess-and-verify** method to pin down the functional form of the value function. Guess that the value function takes the form

$$v(x,s) = -\frac{1}{\gamma}\frac{1}{A}\exp\left\{-\gamma(Ax+By(s)+D)\right\},$$

where A, B, and D are constants to be determined.

It is convenient to fix $A = R/(R-1)$ at the outset; this value will be confirmed below by matching coefficients on the FOC.

By the envelope condition, we have

$$v'(x,s) = u'(c) \implies c(x,s) = \frac{R-1}{R}x + By(s) + D.$$

By definition $a(x,s) = x - c(x,s)$, we have

$$a(x,s) = x - \left(\frac{R-1}{R}x + By(s) + D\right) = \frac{1}{R}x - By(s) - D.$$

The FOC for the Bellman equation is given by

$$u'(c) = \beta R\mathbb{E}[v'(y(s')+Ra',s')|s].$$

By matching the coefficients on both sides, we can solve for A, B and D. And finally we have[6]

$$c(x,s) = \frac{R-1}{R}\left\{x + \frac{\psi}{R-\psi}y(s) + \frac{R(1-\psi)}{(R-\psi)(R-1)}\bar{y}\right\} - \frac{\ln\beta R}{\gamma(R-1)} - \frac{\gamma(R-1)}{(R-\psi)^2}\frac{\sigma^2}{2},$$

where the first term is the permanent value of total resources by the permanent income hypothesis[7], the second term measures the relative impatience[8], and the third term is the

---

[6] Mathematical fact (expectation of the exponential of a normally distributed random variable): If $\varepsilon \sim \mathcal{N}(0,\sigma^2)$, then for any constant t, we have $\mathbb{E}[\exp(t\varepsilon)] = \exp(1/2t^2\sigma^2)$.

[7] The *Permanent Income Hypothesis (PIH)*, developed by Milton Friedman, states that consumption is determined by permanent (long-run expected) income rather than current income. The *permanent income* is the annuity value of total lifetime wealth. In our context, the "permanent value of total resources" $\frac{R-1}{R}\left\{x + \frac{\psi}{R-\psi}y(s) + \frac{R(1-\psi)}{(R-\psi)(R-1)}\bar{y}\right\}$ represents this annuity value: it captures how much the agent can sustainably consume each period given their current assets x, current endowment shock $y(s)$, and expected future endowments. The agent consumes this permanent income, supplemented by precautionary savings motives. The permanent income framework is useful because it separates consumption into a predictable component (based on permanent income) and a precautionary component (due to income uncertainty).

[8] When $\beta R < 1$, the agent is "impatient" in the sense that they value consumption today more than the investment opportunity. This term is independent of the state s, so it acts as a constant adjustment to consumption. The term is positive, and higher impatience (lower β) increases consumption relative to permanent income. Conversely, a higher return R reduces the magnitude of this adjustment.

precautionary saving motive[9].

Moreover, we are interested in the policy function for the "cash-in-hand" variable x, namely the evolution from x to x'. By definition, we have

$$\begin{cases} x = y(s) + Ra \\ x' = y(s') + Ra' \\ c = y(s) + Ra - a' = x - a' \end{cases}$$

These imply that

$$\begin{aligned} x' &= y(s') + R(x - c) \\ &= R\left\{ x - \underbrace{\frac{R-1}{R}\left\{ x + \frac{\psi}{R-\psi}y(s) + \frac{R(1-\psi)}{(R-\psi)(R-1)}\bar{y} - \frac{\ln\beta R}{\gamma(R-1)} - \frac{\gamma(R-1)}{(R-\psi)^2}\frac{\sigma^2}{2} \right\}}_{c} \right\} \\ &\quad + \underbrace{\psi y(s) + (1-\psi)\bar{y} + \varepsilon}_{y(s')} \\ &= x + \left\{ \frac{\psi(1-\psi)}{R-\psi}(y(s) - \bar{y}) + \varepsilon \right\} + \frac{\gamma(\beta-1)}{(R-\psi)^2}\frac{\sigma^2}{2} + \frac{\ln\beta R}{\gamma(R-1)}. \end{aligned}$$

We can see that the second term is negative if $y(s) < \bar{y}$, which means that the agent will borrow more to smooth consumption when the current endowment is low, and save more when the current endowment is high. The third term measure the risk variance, which indicates that the agent will save more to self-insure against the risk when the income is more volatile. The last term measures the relative impatience, which indicates that the agent will save more when they are more patient.

2.3 General Equilibrium

2.3.1 Determining the Equilibrium Interest Rate R

In the partial equilibrium case, we take the interest rate R as given. However, in general equilibrium, the interest rate is endogenously determined by the market clearing condition. Specifically, the interest rate must adjust to ensure that the total demand for assets equals the total supply of assets in the economy:

$$\sum_i a_i = 0.$$

The other way to interpret this market clearing condition is that the total savings of the agents must equal the total borrowing of the agents.

---

[9] The third term $-\frac{\gamma(R-1)}{(R-\psi)^2}\frac{\sigma^2}{2}$ represents the reduction in consumption due to income uncertainty, which motivates precautionary saving. This term is proportional to σ^2 (the variance of the income shock), γ (the coefficient of absolute risk aversion), and depends on the persistence parameter ψ. When income is more uncertain (larger σ^2) or the agent is more risk-averse (larger γ), this term becomes more negative, reducing consumption and increasing savings. This is the "precautionary saving motive": agents save more to build a buffer against future income shocks. The precautionary saving effect is a key mechanism in incomplete-market models that generates substantial asset accumulation even when agents are impatient.

Recall that we construct the cash-in-hands variable $x = y(s) + Ra$. The market clearing condition implies that the total cash-in-hands must equal the total endowment in the economy:

$$\sum_i x_i = \sum_i y(s) + R \sum_i a_i = \sum_i y(s) = \bar{y}.$$

The last equality is true since we assume there is a continuum of agents with measure 1.

From the previous section, we have the policy function for x:

$$x' = x + \left\{ \frac{\psi(1-\psi)}{R-\psi}(y(s) - \bar{y}) + \varepsilon \right\} + \frac{\gamma(\beta-1)}{(R-\psi)^2}\frac{\sigma^2}{2} + \frac{\ln \beta R}{\gamma(R-1)}.$$

The market clearing condition then implies that

$$\begin{aligned} \int x_i' &= \int x_i + \frac{\psi(1-\psi)}{R-\psi}\left(\int y_i - \bar{y}\right) + \int \varepsilon_i + \frac{\gamma(\beta-1)}{(R-\psi)^2}\frac{\sigma^2}{2} + \frac{\ln \beta R}{\gamma(R-1)} \\ \Longrightarrow 0 &= 0 + \frac{\psi(1-\psi)}{R-\psi}(\bar{y} - \bar{y}) + \frac{\gamma(\beta-1)}{(R-\psi)^2}\frac{\sigma^2}{2} + \frac{\ln \beta R}{\gamma(R-1)} \end{aligned}$$

In order for the market to clear, we must have

$$\underbrace{\frac{\gamma(\beta-1)}{(R-\psi)^2}\frac{\sigma^2}{2} + \frac{\ln \beta R}{\gamma(R-1)}}_{\text{Drift}} = 0.$$

Thus, R has to be such that

$$\frac{\ln \beta R}{\gamma(R-1)} = -\frac{\gamma(\beta-1)}{(R-\psi)^2}\frac{\sigma^2}{2}.$$

If we consider the $R-\beta$ plane, the above equation pins down the downward-sloping relationship between R and β:

$$\beta = \frac{1}{R}\exp\left\{ -\frac{\gamma^2(R-1)^2}{(R-\psi)^2}\frac{\sigma^2}{2} \right\}.$$

Note that here the equality does not imply any causal relationship. β is exogenously given, and R is endogenously determined by the market clearing condition.

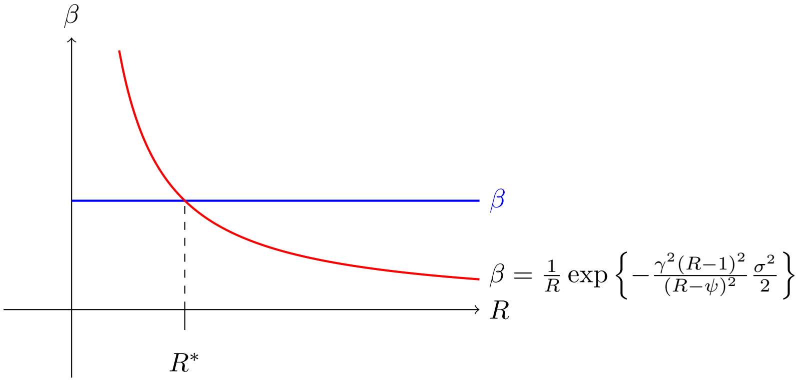

From the graphical intuition, we can see that the general equilibrium interest rate R^* increases when there is

- less volatility in the endowment process (smaller σ^2),
- less risk aversion (smaller γ), and
- less persistence in the endowment process (smaller ψ).

2.3.2 Permanent Income Hypothesis

In addition, we assume that ψ in the AR(1) process is 0. That is, the endowment process is i.i.d.

Then the policy function for x can be rewritten as

$$\begin{aligned} c(x,s) &= \frac{R-1}{R}\left[x+\frac{\bar{y}}{R-1}\right] \\ &= \frac{R-1}{R}x+\frac{\bar{y}}{R} \\ &= \frac{1}{1+r}(rx+\bar{y}). \end{aligned}$$

Note that here we do not have the drift term by the previous argument under general equilibrium. The last equality holds because $R = 1+r$. $c(x,s) = \frac{1}{1+r}(rx+\bar{y})$ has economic intuition: the agent consumes present value of the average income $\bar{y}$ and the return on current cash-in-hand x. This is the permanent income hypothesis: the agent consumes a constant fraction of their permanent income in each period and each state.

The policy function for x can be rewritten as

$$x' = x + \varepsilon(s').$$

2.3.3 Restrictive Borrowing Constraint

In the model assumption, the borrowing constraint could be more restrictive than the natural borrowing limit. Here we assume

- $a \geq \phi$ where $\phi < 0$.
- $y(s)$ is i.i.d. over all states.

The Bellman equation is given by

$$\begin{aligned} v(x) = \max_{a'} u(x-a') + \beta \sum_{s'} \pi(s') v(y(s') + Ra') \\ \text{s.t. } a' \geq \phi. \end{aligned}$$

Define

$$\begin{aligned} \hat{a} &= a - \phi, \\ \tilde{y}(s) &= y(s) + (R-1)\phi. \end{aligned}$$

Then the budget constraint is

$$c = x - a' = (x-\phi) - (a'-\phi) := z - a,$$

where we define $z = x - \phi$.

Moreover, we show that z can be interpreted as the cash-in-hand variable under the restrictive borrowing constraint. Specifically, we have

$$\begin{aligned} R\hat{a} + \tilde{y}(s) &= R(a - \phi) + y(s) + (R - 1)\phi \\ &= Ra - R\phi + y(s) + R\phi - \phi \\ &= Ra + y(s) - \phi \\ &= x - \phi \\ &= z. \end{aligned}$$

Remark (Motivation and Economic Intuition for the Change of Variables).

> The primary goal of introducing these new variables is to mathematically transform a model with a negative borrowing limit ($a \geq \phi$, where $\phi < 0$) into an equivalent, highly standard model with a *strict zero-borrowing limit* ($\hat{a} \geq 0$).
>
> - $\hat{a} = a - \phi$ **(Wealth Buffer):** Since ϕ is the absolute bankruptcy line, $\hat{a}$ measures the agent's "distance to default." If the agent borrows to the maximum limit ($a = \phi$), their buffer is exactly $\hat{a} = 0$.
> - $\tilde{y}(s) = y(s) + (R-1)\phi$ **(Net Disposable Income):** Note that $(R-1)\phi = r\phi$ represents the perpetual interest payment required to service the maximum possible debt. Thus, $\tilde{y}(s)$ is the agent's effective income *after* deducting this worst-case mandatory interest payment.
> - $z = x - \phi$ **(Effective Cash-in-Hand):** This measures how far the agent's current total resources are above the absolute minimum threshold of survival.
>
> Normalizing the constraint to $\hat{a} \geq 0$ significantly simplifies the analytical proofs and the numerical computation of the Bellman equation. It also bounds the ergodic distribution in the phase diagrams below: working with z and $\hat{a}$ cleanly anchors the boundaries of the state space, letting us identify the lower bound ($z_{\min}$, where the poorest agent is constrained) and the upper bound ($z_{\max}$), and so establish that the stationary distribution of assets has compact support.

So the Bellman equation can be rewritten as

$$\begin{aligned} &\hat{v}(z) = \max_{\hat{a}'} u(z - \hat{a}') + \beta \sum_{s'} \pi(s')\hat{v}(R\hat{a}' + \tilde{y}(s')) \\ &\text{s.t. } \hat{a}' \geq 0. \end{aligned}$$

Similarly, the FOC is given by

$$u'(c(z)) \geq \beta R \sum_{s'} \pi(s')\hat{v}'(R\hat{a}' + \tilde{y}(s')).$$

We are interested in how the policy functions $c(z)$ and $\hat{a}(z)$ behave with respect to z.

Let $\hat{z}$ be such that the borrowing constraint binds, i.e.,

$$a' = \phi \iff \hat{a}' = 0.$$

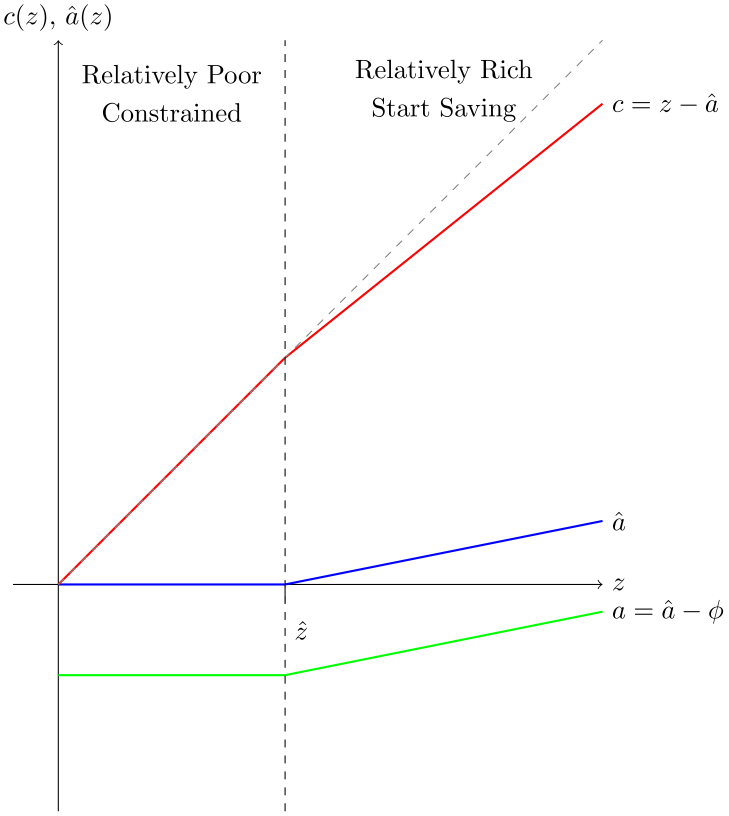

Thus, this allows us to have a general picture of the ergodic distribution of assets in the economy.[10]

---

[10] In the context of heterogeneous agent models, an "ergodic distribution" (often used interchangeably with "stationary distribution") is the unique, long-run probability distribution of agents across all possible states (e.g., wealth and income) that the economy eventually settles into, *regardless of its initial starting point.*

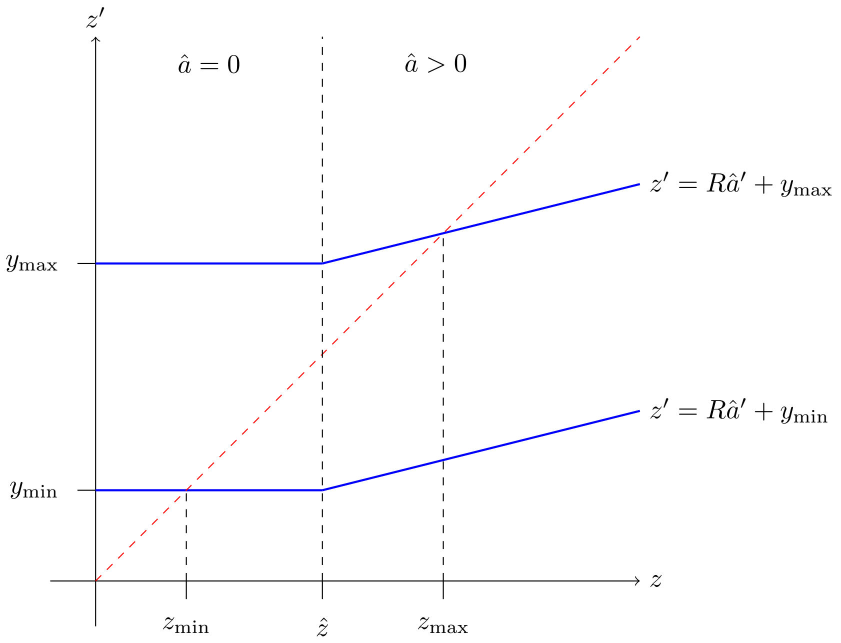

2.3.4 Compute the Stationary Distribution of Assets

Up to this point, we have treated the asset choice a as a continuous variable to derive elegant analytical results (like the Euler equation and supermartingale convergence). However, computers cannot solve Bellman equations over a continuous, infinite state space. To compute the model numerically, we must *discretize* the state space. We assume:

- **Income Grid:** The endowment shock $y(s)$ takes finite discrete values $y(s) \in \{y_1, y_2, \ldots, y_M\}$.
- **Asset Grid:** We define a finite grid of possible asset holdings $\Omega = \{a_1, a_2, \ldots, a_N\}$, where a_1 is typically the borrowing limit ($a_1 = \phi$) and a_N is an upper bound chosen large enough so that it is rarely binding.

By restricting the agent to only choose tomorrow's assets from this predefined grid, the Bellman equation for the agent's problem is given by

$$v(a, s) = \max_{a' \in \Omega} \left[u(Ra + y(s) - a') + \beta \sum_{s'} \pi(s'|s) v(a', s') \right].$$

Because the choice set is finite, we do not need to rely on first-order conditions. The computer simply evaluates the right-hand side for every possible $a' \in \Omega$ and picks the maximum. From this, we obtain a discrete policy function:

$$a' = g(a, s) \in \Omega.$$

Now we keep track of the unconditional distribution of (a, s), denoted by $\lambda(a, s)$:

$$\begin{aligned}
&\lambda_{t+1}(a', s') \\
=&\Pr(a_{t+1} = a', s_{t+1} = s') \\
=&\sum_{s_t}\sum_{a_t} \Pr(a_{t+1} = a' | a_t = a, s_t = s) \Pr(s_{t+1} = s' | s_t = s) \Pr(a_t = a, s_t = s) \\
=&\sum_{s} \sum_{a:a'=g(a,s)} \pi(s'|s)\lambda_t(a, s).
\end{aligned}$$

Thus,

$$\lambda_{t+1}(a', s') = \sum_{s} \sum_{a:a'=g(a,s)} \pi(s'|s)\lambda_t(a, s).$$

Denote $\lambda(a, s)$ as the stationary distribution of (a, s) such that

$$\lambda_t(a, s) = \lambda_{t+1}(a, s) = \lambda(a, s), \quad \forall t.$$

The stationary distribution $\lambda(a, s)$ can be interpreted in two equivalent ways:

- the fraction of time an infinitely-lived agent spends in state (a, s).
- the fraction of agents in the economy that are in state (a, s) at any point in time.

Note that under the stationary distribution, the cross-section distribution of agents over (a, s) remains constant over time. But the state of each individual agent, (a, s), can still change over time.

Aggregate average assets (the aggregate net demand for savings) can be computed as

$$\mathbb{E}[a(R)] = \sum_{a}\sum_{s} g(a, s)\lambda(a, s).$$

Remark.

- You might wonder why we don't simply write this as $\sum_a \sum_s a\lambda(a, s)$, which measures the aggregate assets *currently* held today. Mathematically, in a stationary equilibrium, the aggregate assets today must equal the aggregate assets chosen for tomorrow. Thus, $\sum_a \sum_s a\lambda(a, s) \equiv \sum_a \sum_s g(a, s)\lambda(a, s)$. They yield the exact same number.

  Economically, however, we use the policy function $g(a, s)$ because market clearing is a condition on the **active demand for assets**. We want to aggregate agents' optimal saving *decisions* for the next period to see if the loan market clears.

- The interest rate R dictates the agent's optimal policy rule $g(a, s; R)$, which in turn shapes the ergodic distribution $\lambda(a, s; R)$. Therefore, the aggregate savings demand is a highly nonlinear function of R. In the Huggett model, our ultimate goal is to find the equilibrium root R^* such that the excess demand function evaluates to zero: $\mathbb{E}[a(R^*)] = 0$.

Definition 2.2: Stationary Equilibrium in Huggett Model

A *stationary equilibrium* consists of

- interest rate: R,
- policy function: $a' = g(a,s)$,
- stationary distribution: $\lambda(a,s)$, and
- given borrowing limits

such that

- Given R, $g(a,s)$ and $v(a,s)$ solve the agent's problem.
- $\lambda(a,s)$ is the stationary distribution of (a,s) given $g(a,s)$ and $\pi(s'|s)$.
- Given $\lambda(a,s)$, $g(a,s)$ and R, the loan market clears:

$$\sum_a \sum_s g(a,s)\lambda(a,s) = 0$$

Algorithm for Huggett Model Solution

1. Guess an interest rate R.
2. Solve the agent's problem to obtain the policy function $a' = g(a,s)$.
3. Compute the stationary distribution.
    - Start from any λ_0.
    - Iterate $\lambda_{t+1}(a',s') = \sum_s \sum_{a:a'=g(a,s)} \pi(s'|s)\lambda_t(a,s)$ until convergence.
4. Check if the loan market clears:

$$\sum_a \sum_s g(a,s)\lambda(a,s) = \varepsilon.$$

If $\varepsilon > 0$, this means there is too much assets (and thus savings) in the economy, which implies that R is too high. Hence, we need to decrease R and go back to step 2. If $\varepsilon < 0$, this means there is too much borrowing in the economy, which implies that R is too low. Hence, we need to increase R and go back to step 2. If $\varepsilon = 0$, then we have found the equilibrium interest rate.

2.4 Aiyagari Model

Motivation: From Endowment to Production

The Huggett model is an ***endowment economy*** *where the only asset is a risk-free bond in zero net supply (for every borrower, there must be a saver). The Aiyagari model takes* *a massive step forward by embedding heterogeneous agents into a Neoclassical* ***production economy****. Specifically, agents no longer just trade IOUs with each other. They save by accumulating physical capital ($k \geq 0$). Thus, the aggregate savings in the economy correspond to the aggregate physical capital stock ($K > 0$), which is used for production.*

Assume:

- Incomplete market: agents can save only via accumulating physical capital k.
- Production: $y = F(K, L)$, typically a constant returns to scale (CRS) technology.
- Shocks: idiosyncratic i.i.d. (or Markov) shocks to an individual's labor productivity s.
- Inelastic labor supply: $L = 1$.
- Law of motion for capital:

  $$k(s^t) = (1-\delta)k(s^{t-1}) + i(s^{t+1}),$$

  where $\delta \in [0, 1]$ is the depreciation rate of capital.
- Budget constraint:

  $$c(s^t) + i(s^t) \leq wl(s^t) + rk(s^{t-1}),$$

  where w is the wage rate, r is the rental rate of capital, $rk(s^{t-1})$ is the rent earned from the $(t-1)$-capital, and $wl(s^t)$ is the labor income at time t.

Remark (Relation to the Huggett Model).

> We can map the Aiyagari budget constraint back to the familiar Huggett format. By substituting the investment $i(s^t) = k(s^t) - (1-\delta)k(s^{t-1})$ into the budget constraint, we get:
>
> $$c(s^t) + k(s^t) - (1-\delta)k(s^{t-1}) \leq wl(s^t) + r_k k(s^{t-1}).$$
>
> Rearranging the terms to group the current and future capital yields:
>
> $$c(s^t) + k(s^t) \leq wl(s^t) + (1 + r_k - \delta)k(s^{t-1}).$$
>
> This matches the Huggett budget constraint $c + a' \leq y(s) + Ra$, where:
>
> - $k(s^t)$ acts as the new asset a'.
> - $wl(s^t)$ acts as the stochastic income $y(s)$.
> - $(1 + r_k - \delta)$ acts as the gross risk-free interest rate R.

We define the aggregate variables as

$$K = \sum_k \sum_s \lambda(k, s) g(k, s),$$

$$L = \sum_k \sum_s \lambda(k, s) l(s) = 1.$$

For the firm, the static profit maximization problem is given by:

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

The FOCs yield the wage rate and the rental rate of capital:

$$\begin{aligned} r_k &= F_K(K,L), \\ w &= F_L(K,L). \end{aligned}$$

However for the household, we consider the interest rate to be $r = F_K(K,L) + 1 - \delta$. Note that it is crucial to distinguish between the *rental rate of capital* (r_k) paid by the firm and the *gross return* (R) received by the household. The firm rents capital from the household, uses it for production, and pays the rental rate $r_k = F_K$. During production, the capital depreciates by δ. The household then takes back the undepreciated capital $(1-\delta)K$. Therefore, the gross return to the household for saving one unit of capital is the rental payment plus the undepreciated principal: $r = r_k + 1 - \delta = F_K(K,L) + 1 - \delta$.

The agent's problem is then given by

$$\begin{aligned} v(k,s) = \max_{c,k'} \;& u(c) + \beta \sum_{s'} \pi(s'|s) v(k',s') \\ \text{s.t. } & c + k' \le (1 + F_K - \delta)k + wl(s) \\ & k' \ge 0. \end{aligned}$$

The diagram below summarizes the joint determination of aggregate capital (K) and the interest rate (r) across three benchmark macroeconomic models, by plotting the firm's *aggregate capital demand* (blue curve) against the household sector's *aggregate asset supply* (red curve).

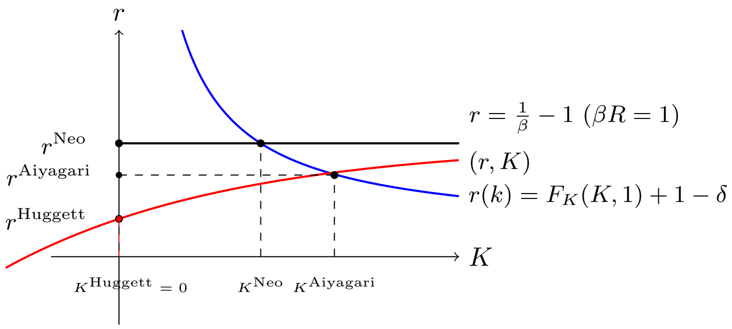

- **The Blue Curve (Capital Demand):** This represents the firm's first-order condition, $r = F_K(K,1) - \delta$. It is downward-sloping due to the diminishing marginal product of capital.

- **The Neo-classical Case (Complete Markets Benchmark):** In a standard representative-agent model (or with complete markets), there is no uninsurable idiosyncratic risk. The steady-state Euler equation is simply $1 = \beta(1+r)$, which completely and permanently pins down the interest rate at $r^{\text{Neo}} = \frac{1}{\beta} - 1$. Because the agent has no precautionary saving motive, the long-run asset supply curve is perfectly elastic (the horizontal black line).

The equilibrium capital K^{Neo} is strictly determined by where the firm's downward-sloping demand intersects this horizontal line.

- **The Red Curve (Incomplete Markets Asset Supply):** In the presence of uninsurable idiosyncratic risk, agents accumulate precautionary savings. As r increases, the incentive to save grows, making the supply curve upward-sloping. Crucially, as $r \to \frac{1}{\beta} - 1$, agents become infinitely patient relative to the market return, and their asset demand diverges to infinity (as proven earlier via the supermartingale theorem). Thus, the red curve stays strictly below the black line and asymptotes to it.

- **The Huggett Case (Pure Exchange Economy):** In Huggett, there is no physical capital, meaning the net supply of assets must clear at exactly zero ($K = 0$). The equilibrium interest rate r^{Huggett} is found where the red supply curve intersects the y-axis. Notice that $r^{\text{Huggett}} < r^{\text{Neo}}$: to discourage agents from over-saving in a zero-net-supply economy, the market interest rate must be severely depressed.

- **The Aiyagari Case (Production + Incomplete Markets):** This is the grand synthesis. General equilibrium is reached at the intersection of the red savings supply curve and the blue capital demand curve.

- **Key Takeaway (Aiyagari vs. Neo-classical):** Because the red supply curve strictly lies below the black line (due to precautionary savings), the Aiyagari intersection must occur further down the blue demand curve. Consequently, $K^{\text{Aiyagari}} > K^{\text{Neo}}$ and $r^{\text{Aiyagari}} < r^{\text{Neo}}$. The idiosyncratic risk forces agents to over-accumulate capital for self-insurance, which in turn drives down the equilibrium interest rate compared to the frictionless benchmark.

Definition 2.3: Stationary Equilibrium in Aiyagari Model

A *stationary equilibrium* consists of

- stationary distribution: $\lambda(k, s)$,
- value function: $v(k, s)$,
- policy function: $k' = g(k, s)$,
- prices: r and $w(r)$,
- aggregate capital: K

such that

- Prices r and $w(r)$ satisfy:

$$\begin{aligned} w &= F_L(K, 1), \\ r &= F_K(K, 1) + 1 - \delta. \end{aligned}$$

- $v(k, s)$ and $g(k, s)$ solve the agent's problem given r and w.
- $\lambda(k, s)$ is the stationary distribution of (k, s) given $g(k, s)$ and $\pi(s'|s)$:

$$\lambda(k', s') = \sum_s \sum_{k:k'=g(k,s)} \pi(s'|s)\lambda(k, s).$$

- The stationary distribution $\lambda(k, s)$ and policy function $g(k, s)$ generate aggregate capital K:

$$\sum_k \sum_s g(k, s)\lambda(k, s) = K.$$

Algorithm for Aiyagari Model Solution

1. Guess an aggregate capital K_0.
2. Obtain the prices r and w from the firm's first-order conditions:

$$\begin{aligned} r &= F_K(K_0, 1) + 1 - \delta, \\ w &= F_L(K_0, 1). \end{aligned}$$

3. Solve the agent's recursive problem to obtain the policy function $k' = g(k, s)$.
4. Obtain the stationary distribution $\lambda(k, s)$.
5. Update K_0 by K_1[a]:

$$K_1 = \sum_k \sum_s g(k, s)\lambda(k, s).$$

6. Iterate until $K_0 = K_1$.

---

[a] A more robust way to update K is to use a convex combination of K_0 and K_1: $K^* = \varepsilon K_0 + (1-\varepsilon)K_1$, where $\varepsilon \in (0,1)$ is called the relaxation parameter.

2.5 Krusell-Smith Model

Both the Huggett and Aiyagari models are steady-state models. While individuals experience idiosyncratic shocks and their personal wealth fluctuates, the aggregate economy never changes. There are no business cycles, and the wealth distribution $\lambda(k, s)$ is stationary; prices (r and w) remain constant forever. The **?** *model breaks this tranquility by introducing* ***aggregate shocks*** *(e.g., TFP shocks, z_t). This seemingly innocent addition creates a serious technical challenge:*

- ***Fluctuating Prices:*** *Because TFP fluctuates, aggregate capital K_t and labor demand fluctuate, meaning r_t and w_t now change over time.*
- ***The Curse of Dimensionality:*** *To make an optimal saving decision today, an agent must forecast tomorrow's prices (r_{t+1}, w_{t+1}). To forecast tomorrow's prices, they must forecast tomorrow's aggregate capital K_{t+1}. But K_{t+1} depends on how everyone in the economy saves today. Therefore, the agent must know the* ***entire current wealth distribution*** *λ_t to predict the future.*
- ***The Infinite-Dimensional State Space:*** *The distribution λ_t is an infinite-dimensional object. Putting an infinite-dimensional object into a Bellman equation as a state variable makes it computationally impossible to solve using traditional grid methods.*

Assume:

- Aggregate shock z_t to productivity such that

$$y_t = z_t F(K_t, L_t).$$

  - Assume z_t follows a finite-state Markov process.
  - z_t can be interpreted as an aggregate technology (business cycle) shock.
- Idiosyncratic shock s_t to labor productivity (also a Markov process).

Therefore, the agents' wealth will depend on both the aggregate shock z_t and the idiosyncratic shock s_t. Pay special attention to the fact that the aggregate distribution λ_t will now vary over time depending on the realization of z_t.

The recursive problem of the agent is given by

$$\begin{aligned} v(k, s, \lambda, z) &= \max_{c, k'} u(c) + \beta \mathbb{E}\left[v(k', s', \lambda', z')|s, \lambda, z\right] \\ \text{s.t. } & c + k' \leq (r(\lambda, z) + 1 - \delta)\, k + w(\lambda, z) l(s), \\ & \lambda' = H(\lambda, z, z'). \end{aligned}$$

where λ is the distribution of capital across agents, and $H(\cdot)$ is the highly complex aggregate law of motion for the distribution of capital. Notice that prices r and w are now explicitly functions of the aggregate state (λ, z).

From the Bellman equation, we can obtain the policy function for capital $k' = g(k, s, \lambda, z)$ by combining the budget constraint and the law of motion for capital.

In the aggregate, the market clearing condition is

$$K_t = \int k\lambda_t(k, s)\mathrm{d}k\mathrm{d}s.$$

Definition 2.4: Recursive Competitive Equilibrium

A *recursive competitive equilibrium* consists of

- value function $v(k, s, \lambda, z)$,
- policy function $k' = g(k, s, \lambda, z)$,
- pricing functions $r(\lambda, z)$ and $w(\lambda, z)$, and
- Aggregate law of motion H that maps (λ, z, z') to λ'

such that

- **Individual Optimization:** Given H, r, and w, the functions v and g solve the agent's Bellman equation.
- **Firm Optimization:** Prices satisfy the firm's first-order conditions: $r(\lambda, z) = zF_K(K, L) - \delta$ and $w(\lambda, z) = zF_L(K, L)$.
- **Market Clearing:** Aggregate capital equals the sum of all individual asset holdings:
$$K = \int k\mathrm{d}\lambda(k, s).$$
- **Consistency (Rational Expectations):** The perceived aggregate law of motion H is consistent with the actual law of motion generated by the aggregation of individual policy functions g:
$$\lambda'(k', s') = \int \int_{k'=g(k,s,\lambda,z)} \pi(s'|s)\pi(z'|z)\mathrm{d}\lambda(k, s).$$

Algorithm for Krusell-Smith Model Solution

1. Characterize $\lambda(k, s)$ by a finite number of moments: $m = \{m_1, m_2, \cdots, m_I\}$.
2. Assume a functional form for H that maps m to m': $m' = H(m, z, z')$.
3. Guess the m's and H.
4. Solve the Bellman equation to obtain the value function $v(k, s, m, z)$ and the policy function $k' = g(k, s, m, z)$.
5. Simulate the economy:

- Draw realizations for $\{z_t\}_{t=0}^T$.
- Simulate the paths of $\{s_t\}_{t=0}^T$ for a large number of agents.
- Obtain the simulated paths of $\{k_t\}_{t=0}^T$ for each agent by using the policy function $g(k, s, m, z)$.
- Assemble the simulated paths of k_t to obtain the simulated paths of m_t.

6. Update H based on the simulated paths of m_t.
7. Iterate until convergence.

It turns out that tracking only the first moment ($m_1 = K$) is essentially sufficient: the R^2 of the forecasting regression $\ln K' = a_z + b_z \ln K$ is typically above 0.999. Why does this "approximate aggregation" work so well?

The intuition lies in the shape of the policy functions. The poorest agents face binding borrowing constraints, so their saving policy is highly non-linear. But the poorest agents hold almost zero capital, so their non-linear behavior barely moves the macroeconomic aggregate K. The richest agents, by contrast, hold the vast majority of the economy's capital, and for them the precautionary motive is negligible, making the policy function $k' = g(k, \dots)$ nearly linear in k.

Since the agents who actually matter for the aggregate sum behave linearly, aggregate capital tomorrow (K') depends mostly on aggregate capital today (K), with relatively little dependence on how that capital is distributed among the rich. The infinite-dimensional distribution λ therefore collapses, to a very good approximation, into a single number K.

Remark (Chapter Summary).

> - **Single-asset incomplete markets break perfect risk sharing.** Restricting trade to a risk-free bond exposes households to the variance of $y(s)$. With $u''' > 0$, this generates a precautionary saving motive absent from the complete-markets benchmark.
> - **Cash-in-hand $x = y + Ra$ is the natural state under i.i.d. shocks.** It collapses (a, s) to a single sufficient statistic. Persistence (e.g., AR(1)) breaks the collapse and forces a two-dimensional state (x, s).
> - **The Huggett model.** Continuum of households, single risk-free bond in zero net supply. Equilibrium r pins down the cross-sectional distribution $\lambda(a, s)$ as the fixed point of the policy-induced transition. Computation: outer loop on r, inner VFI plus iteration on λ.
> - **The Aiyagari model.** Adds production. Aggregate household savings now equal aggregate physical capital. Precautionary motives drive $r^* < 1/\beta - 1$ strictly. The equilibrium picture is the intersection of an upward-sloping household-supply curve with a downward-sloping firm-demand curve.
> - **Krusell–Smith adds aggregate uncertainty.** The wealth distribution becomes a dynamic state, in principle infinite-dimensional. Approximate aggregation: a single moment (K) suffices because the rich behave nearly linearly.

- **Chapter 11 extends the algorithmic story.** The same three-loop computation reappears in the Aiyagari computation chapter, with more numerical detail (EGM, Howard improvement, sparse eigenvector solves).

Part II

Growth, Business Cycles, and Quantitative Macroeconomics

Lectures by Kai-Jie Wu

Part III

Problem Sets and Solutions

Subject Index

Author Index