Advanced Macroeconomics Theory

Maria-Jose Carreras-Valle & Kai-Jie Wu LaTeX by Rui Zhou
Penn State University

Spring 2026

Contents

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; $\pi(s^\tau \mid s^t)$ 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 1

Complete Markets under Uncertainty

Remark (Notation in This Chapter).

Symbol	Meaning
$i \in \{1, \ldots, I\}$	Index for the I heterogeneous agents
λ^i	Pareto weight on agent i in the planner's problem
μ^i	Lagrange multiplier on agent i's lifetime budget constraint
$\theta(s^t)$	Multiplier on the resource constraint at history s^t
$a^i(s_{t+1}\|s^t)$	Agent i's holding of the one-period-ahead Arrow security
$A^i(s^t)$	natural debt limit (PV of future endowment stream)
$Q(s_{t+1}\|s^t)$	Pricing kernel in sequential trading
$V^i(a, s)$	Recursive value function
$P(v)$	Recursive Pareto frontier (max utility for agent 2 given promised utility v to agent 1)
v, w_s	Promised utility today and continuation promise in state s
m_{t+1}	stochastic discount factor (Asset Pricing remark)

By "uncertainty" we mean that the future is not deterministic: multiple states of the world are possible, and agents' endowments differ across states. "Complete markets" means that the menu of available financial assets is rich enough to span every state—agents can purchase state-contingent claims that pay off in any specific future contingency, and therefore can smooth consumption perfectly across both states and time. In an incomplete market, by contrast, the asset menu is not rich enough to insure against every shock, and consumption smoothing is necessarily imperfect.

Most of this material was covered last semester. We revisit it here as the foundation for the incomplete-markets analysis introduced in subsequent chapters. The central question this chapter prepares us to answer is: ***in an endowment economy with uncertainty and heterogeneous agents, how does market incompleteness alter the equilibrium relative to the complete-markets benchmark?***

Assume:

- Stochastic event: for $t \geq 0$, $s_t \in S$.
- History of events: $s^t = \{s_0, s_1, \cdots, s_t\}$, which is public information.
- Probability of s^t:
  - $\pi(s^t)$: unconditional probability.
  - $\pi(s^\tau | s^t)$: conditional probability of s^τ given history s^t for $\tau > t$.
- I agents: $i = 1, \cdots, I$.
- Stochastic endowments: $y^i(s^t)$, which is non-storable.
- Consumption allocations: $c^i(s^t)$ such that $c^i(s^t) \geq 0$, $\forall i$, $\forall s^t$.
- Preferences:
$$u(c^i) = \sum_t \sum_{s^t} \beta^t \pi(s^t) u(c^i(s^t)),$$
where u is assumed to be concave, strictly increasing, differentiable, and satisfy Inada condition: $\lim_{c \to 0} u'(c) = \infty$.

It is helpful to visualize the structure of histories before going further. With $S = \{s^a, s^b\}$ binary and three periods, the set of histories s^t is a binary tree whose nodes are labeled by the cumulative draws to date:

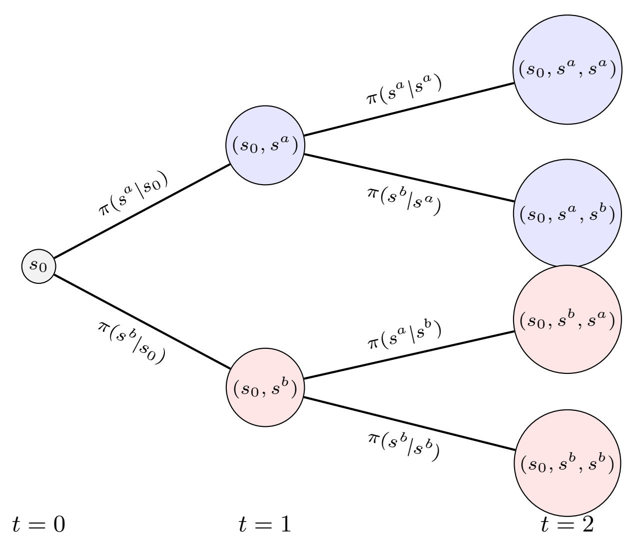

A few features deserve emphasis:

- Each **node** is a history s^t, not just a state s_t. Two nodes can have the same current s_t but represent different histories—e.g., (s_0, s^a, s^b) and (s_0, s^b, s^b) both end at s^b at date $t = 2$ but have walked different paths to get there. Allocations $c^i(s^t)$ and asset holdings $a^i(s^t)$ are indexed by the full history, not by s_t.
- In the Arrow–Debreu formulation (next section), all trades occur at the root node s_0. A state-contingent claim $q(s^t|s_0)$ is the date-0 price of consumption delivered if and only if the realized path through this tree reaches the node s^t.

- In the sequential-trading formulation, the agent stands at some node s^t each period and trades only one-step-ahead claims to the child nodes $s^{t+1} = (s^t, s_{t+1})$.

Definition 1.1: Feasible Allocation

An allocation $\left\{c^i(s^t)\right\}_{i,t,s^t}$ is feasible if

$$\sum_i c^i(s^t) \leq \sum_i y^i(s^t), \quad \forall t, \quad \forall s^t.$$

Definition 1.2: Pareto Optimal

A feasible allocation is Pareto optimal if any other feasible allocation that makes one agent strictly better off must make at least one agent strictly worse off.

1.1 Find Pareto Optimal Allocations: Social Planner

The social planner's problem is given by

$$\begin{aligned} \max_{\{c^i(s^t)\}_{i,t,s^t}} & \sum_i \lambda^i \left(\sum_t \sum_{s^t} \beta^t \pi(s^t) u(c^i(s^t)) \right) \\ \text{s.t.} \ & \sum_i c^i(s^t) \leq \sum_i y^i(s^t), \quad \forall t, \quad \forall s^t \end{aligned}$$

The Lagrange is given by

$$\mathcal{L} = \sum_t \sum_{s^t} \beta^t \pi(s^t) \left(\sum_i \lambda^i u(c^i(s^t)) \right) + \sum_t \sum_{s^t} \theta(s^t) \left(\sum_i y^i(s^t) - \sum_i c^i(s^t) \right)$$

where $\theta(s^t)$ is the Lagrange multiplier for the resource constraint at state s^t.

The first-order condition for $c^i(s^t)$ is given by

$$\frac{\partial \mathcal{L}}{\partial c^i(s^t)} = \lambda^i \beta^t \pi(s^t) u'(c^i(s^t)) - \theta(s^t) = 0 \quad \Longrightarrow \quad \theta(s^t) = \lambda^i \beta^t \pi(s^t) u'(c^i(s^t)).$$

This implies that for any two agents i and j, we have

$$\frac{u'(c^i(s^t))}{u'(c^j(s^t))} = \frac{\lambda^j}{\lambda^i}, \quad \forall t, \quad \forall s^t.$$

This is the condition for Pareto optimality.

Specifically,

$$\frac{u'(c^i(s^t))}{u'(c^1(s^t))} = \frac{\lambda^1}{\lambda^i}, \quad \forall t, \quad \forall s^t.$$

And this gives

$$c^i(s^t) = (u')^{-1} \left(\frac{\lambda^1}{\lambda^i} u'(c^1(s^t)) \right), \quad \forall t, \quad \forall s^t.$$

Feasibility constraint requires that

$$\sum_i c^i(s^t) = \sum_i y^i(s^t), \quad \forall t, \quad \forall s^t.$$

Hence we have

$$\sum_i (u')^{-1}\left(\frac{\lambda^1}{\lambda^i} u'(c^1(s^t))\right) = \sum_i y^i(s^t), \quad \forall t, \quad \forall s^t.$$

This equation pins down $c^1(s^t)$, and thus $c^i(s^t)$ for all i using $c^i(s^t) = (u')^{-1}\left(\frac{\lambda^1}{\lambda^i} u'(c^1(s^t))\right)$.

Remark.

- $c^1(s^t)$ (and thus $c^i(s^t)$ for all i) depends on the Pareto weights $\left\{\lambda^i\right\}_{i=1}^I$ and the total endowments $\sum_i y^i(s^t)$, but not on the distribution of endowments across agents.
- $c^i(s^t)$ does not depend on the actual realization of income (endowment).
- If $\lambda^i = \lambda^j$ for all i, j, then $c^i(s^t) = c^j(s^t)$ for all i, j, and thus $c^i(s^t) = \frac{1}{I}\sum_i y^i(s^t)$, which is the full insurance allocation where all agents have equal share of total endowment in each state.

1.2 Arrow-Debreu Trading (Time 0 Trading)

Suppose that all trades occur at time 0, when agents are going to exchange claims contingent on history s^t at price $q(s^t|s_0)$ (price of consumption at s^t in terms of consumption at s_0). Later we will simply write $q(s^t)$ if there is no ambiguity.

Definition 1.3: Competitive Equilibrium (Time-0 Trading)

A competitive equilibrium consists of

- a set of prices $q(s^t)$, and
- a set of consumption allocations $\left\{c^i(s^t)\right\}_{i,t,s^t}$

such that

- **Agent's Utility Max**:

  Given prices $q(s^t)$, the consumption allocation $\left\{c^i(s^t)\right\}_{t,s^t}$ solves the following problem for each agent i:

$$\begin{aligned} \max_{\{c^i(s^t)\}_{t,s^t}} \quad & \sum_t \sum_{s^t} \beta^t \pi(s^t) u(c^i(s^t)) \\ \text{s.t.} \quad & \sum_t \sum_{s^t} q(s^t) c^i(s^t) \leq \sum_t \sum_{s^t} q(s^t) y^i(s^t), \quad \forall i \end{aligned}$$

- **Market Clearing**:

$$\sum_i c^i(s^t) = \sum_i y^i(s^t), \quad \forall t, \quad \forall s^t.$$

Remark.

Note that in the agent's problem, she faces the "lifetime" budget constraint:

$$\sum_t \sum_{s^t} q(s^t) c^i(s^t) \leq \sum_t \sum_{s^t} q(s^t) y^i(s^t), \quad \forall i.$$

In this problem, this is the **one and only** budget constraint that the agent faces. Take note that this is different from the budget constraints in the sequential problem we will see in the next section, where the agent faces a sequence of period-by-period budget constraints.

For the agent's problem, the Lagrange is given by

$$\begin{aligned} \mathcal{L} &= \sum_t \sum_{s^t} \beta^t \pi(s^t) u(c^i(s^t)) + \mu^i \left(\sum_t \sum_{s^t} q(s^t) y^i(s^t) - \sum_t \sum_{s^t} q(s^t) c^i(s^t) \right) \\ &= \sum_t \sum_{s^t} \beta^t \pi(s^t) u(c^i(s^t)) + \mu^i \sum_t \sum_{s^t} q(s^t) \left(y^i(s^t) - c^i(s^t) \right), \end{aligned}$$

where μ^i is the Lagrange multiplier for the budget constraint.

The first-order condition for $c^i(s^t)$ is given by

$$\frac{\partial \mathcal{L}}{\partial c^i(s^t)} = \beta^t \pi(s^t) u'(c^i(s^t)) - \mu^i q(s^t) = 0 \quad \Longrightarrow \quad \mu^i q(s^t) = \beta^t \pi(s^t) u'(c^i(s^t)).$$

This implies that for any two agents i and j, we have

$$\frac{u'(c^i(s^t))}{u'(c^j(s^t))} = \frac{\mu^i}{\mu^j}, \quad \forall t, \quad \forall s^t.$$

Specifically,

$$\frac{u'(c^i(s^t))}{u'(c^1(s^t))} = \frac{\mu^i}{\mu^1}, \quad \forall t, \quad \forall s^t.$$

This gives

$$c^i(s^t) = (u')^{-1}\left(\frac{\mu^i}{\mu^1}u'(c^1(s^t))\right), \quad \forall t, \quad \forall s^t.$$

From the market clearing condition, we have

$$\sum_i c^i(s^t) = \sum_i (u')^{-1}\left(\frac{\mu^i}{\mu^1}u'(c^1(s^t))\right) = \sum_i y^i(s^t), \quad \forall t, \quad \forall s^t.$$

This allows us to solve for $c^1(s^t)$, and thus $c^i(s^t)$ for all i using $c^i(s^t) = (u')^{-1}\left(\frac{\mu^i}{\mu^1}u'(c^1(s^t))\right)$.

Claim

A competitive equilibrium allocation is Pareto optimal when

$$\mu^i = 1/\lambda^i, \quad \forall i,$$
$$q(s^t) = \theta(s^t), \quad \forall t, \quad \forall s^t.$$

Proof for Claim.

Recall that in the social planner's problem, the first-order condition for $c^i(s^t)$ is given by

$$\lambda^i\beta^t\pi(s^t)u'(c^i(s^t)) - \theta^t(s^t) = 0 \quad \Longrightarrow \quad \theta^t(s^t) = \lambda^i\beta^t\pi(s^t)u'(c^i(s^t)).$$

And this gives

$$\frac{u'(c^i(s^t))}{u'(c^j(s^t))} = \frac{\lambda^j}{\lambda^i}, \quad \forall t, \quad \forall s^t.$$

While in the competitive equilibrium, the first-order condition for $c^i(s^t)$ is given by

$$\beta^t\pi(s^t)u'(c^i(s^t)) - \mu^i q(s^t) = 0 \quad \Longrightarrow \quad \mu^i q(s^t) = \beta^t\pi(s^t)u'(c^i(s^t)).$$

And this gives

$$\frac{u'(c^i(s^t))}{u'(c^j(s^t))} = \frac{\mu^i}{\mu^j}, \quad \forall t, \quad \forall s^t.$$

Hence, if $\mu^i = 1/\lambda^i$ for all i, and $q(s^t) = \theta(s^t)$ for all t and s^t, then the competitive equilibrium allocation is Pareto optimal since their first-order conditions are then identical.

Theorem 1.4: First Welfare Theorem

A competitive equilibrium allocation is Pareto optimal.

Theorem 1.5: Second Welfare Theorem

Any Pareto optimal allocation can be supported by a competitive equilibrium (with transfers of initial endowments).

Algorithm (Solve for Competitive Equilibrium)

1. Fix μ^1. Guess the rest of the μ^i's.
2. Using *FOC* and *feasibility constraint*:

$$\sum_i (u')^{-1}\left(\frac{\mu^i}{\mu^1}u'(c^1(s^t))\right) = \sum_i y^i(s^t), \quad \forall t, \quad \forall s^t,$$

   From this solve for $c^1(s^t)$, and thus $c^i(s^t)$ for all i.

3. Using budget constraint:

$$\sum_t \sum_{s^t} q(s^t)\left(c^i(s^t) - y^i(s^t)\right) = 0, \quad \forall i,$$

   where $q(s^t) = \beta^t \pi(s^t) u'(c^i(s^t))/\mu^i$ from the FOC. (You may plug in any i.)

   Let

$$\sum_t \sum_{s^t} q(s^t)\left(c^i(s^t) - y^i(s^t)\right) = \varepsilon_i.$$

   If the error term $\varepsilon_i > 0$, increase μ^i; if $\varepsilon_i < 0$, decrease μ^i. Iterate until ε_i is close enough to 0 for all i. [a]

4. Update guesses for μ^i's and repeat steps 2 and 3 until ε_i is close enough to 0 for all i.

[a] When $\varepsilon_i > 0$, intuitively it means that the net present value of consumption is greater than the net present value of endowment, which means that the agent is consuming more than what they can afford. So we need to lower c^i. Since $\frac{u'(c^i(s^t))}{u'(c^1(s^t))} = \frac{\mu^i}{\mu^1}$, we need to increase μ^i since $\mu^i \uparrow \implies u'(c^i(s^t)) \uparrow \implies c^i(s^t) \downarrow$.

The economic intuition is that μ^i represents the shadow price of wealth (the Lagrange multiplier on the intertemporal budget constraint) for agent i. When the agent overspends, the penalty for violating the budget must increase. Raising μ^i makes wealth more precious, forcing the agent to value each unit of wealth more highly. This drives up their marginal utility of consumption and forces them to cut back on actual consumption in every state until their lifetime budget is balanced.

Example.

> Suppose there are two agents. There is no aggregate risk in the endowment: $y^1(s^t) + y^2(s^t) = 1$ for all t and s^t. We suppose the two agents have identical preferences u. Solve for the competitive equilibrium.
>
> **Solution.**

FOC gives

$$\frac{u'(c^1(s^t))}{u'(c^2(s^t))} = \frac{\mu^1}{\mu^2}, \quad \forall t, \quad \forall s^t.$$

By the market clearing condition, we have

$$c^1(s^t) + c^2(s^t) = c^1(s^t) + (u')^{-1}\left(\frac{\mu^1}{\mu^2}u'(c^1(s^t))\right) = 1\ (= y^1(s^t) + y^2(s^t)), \quad \forall t, \quad \forall s^t.$$

This implies that $c^1(s^t)$ is constant across s^t: $c^1(s^t) = c^1$ for all s^t. Similarly, $c^2(s^t)$ is also constant across s^t: $c^2(s^t) = c^2$ for all s^t.

FOC also helps pin down the price $q(s^t)$:

$$q(s^t) = \beta^t \pi(s^t) u'(c^1)/\mu^1.$$

Plugging it back into the budget constraint, we have

$$\begin{aligned}
&\sum_t \sum_{s^t} q(s^t)\left(c^i - y^i(s^t)\right) = 0 \\
\implies &\sum_t \sum_{s^t} \frac{\beta^t \pi(s^t) u'(c^i)}{\mu^i}\left(c^i - y^i(s^t)\right) = 0 \\
\implies &\sum_t \sum_{s^t} \beta^t \pi(s^t)(c^i - y^i(s^t)) = 0 \\
\implies &c^i \sum_t \sum_{s^t} \beta^t \pi(s^t) = \sum_t \sum_{s^t} \beta^t \pi(s^t) y^i(s^t) \\
\implies &c^i \cdot \frac{1}{1-\beta} = \sum_t \sum_{s^t} \beta^t \pi(s^t) y^i(s^t) \\
\implies &c^i = (1-\beta)\sum_t \sum_{s^t} \beta^t \pi(s^t) y^i(s^t).
\end{aligned}$$

Note that $\sum_t \sum_{s^t} \beta^t \pi(s^t) y^i(s^t)$ is the net present value of endowments for agent i. So in the equilibrium, agents can be fully insured against the idiosyncratic risk.

Remark.

- **Idiosyncratic vs. Aggregate Risk:** The result that consumption is constant ($c^i(s^t) = c^i$)—i.e., perfect consumption smoothing—relies on the assumption of *no aggregate risk* ($y^1(s^t) + y^2(s^t) = 1$). If the aggregate endowment $Y(s^t) = y^1(s^t)+y^2(s^t)$ were to fluctuate across states, individual consumptions $c^i(s^t)$ would necessarily fluctuate as well. Agents cannot insure away aggregate shocks; they can only share them, so individual consumption would comove positively with the aggregate endowment.

- **Are identical utility functions necessary?** Interestingly, no. Suppose Agent 1 and Agent 2 have completely different utility functions (e.g., $u_1(c) = \ln(c)$ and $u_2(c) = 2\sqrt{c}$). As long as both are risk-averse (strictly concave utility functions), the first-order condition dictates:

$$\frac{u_1'(c^1(s^t))}{u_2'(c^2(s^t))} = \frac{\mu^1}{\mu^2} \quad \text{(constant for all } t, s^t\text{)}$$

Combining this with the market clearing condition $c^1(s^t) + c^2(s^t) = Y$ (where Y is constant without aggregate risk): since both marginal utilities u_1' and u_2' are strictly decreasing, the only way for the ratio to be constant while the sum is constant is for c^1 and c^2 to themselves be constant across all states.

1.3 Sequential Trading

In contrast to Arrow-Debreu trading, where all trades occur at time 0, sequential trading assumes that markets open at every date t and node s^t.

At each node s^t, agents trade a complete set of *one-period-ahead Arrow securities.* Let $a^i(s_{t+1}|s^t)$ denote the quantity of claims owned by agent i at s^t that pays one unit of consumption good tomorrow if and only if state s_{t+1} is realized.[1] $a^i(s^t) > 0$ denotes assets and $a^i(s^t) < 0$ denotes debts. The price of this claim at node s^t is denoted as $Q(s_{t+1}|s^t)$ (often referred to as the *pricing kernel*).

Consequently, the agent faces a sequential budget constraint at every node s^t:

$$c^i(s^t) + \sum_{s_{t+1}|s^t} Q(s_{t+1}|s^t)a^i(s_{t+1}|s^t) \leq y^i(s^t) + a^i(s^t), \quad \forall t, \forall s^t.$$

where $a^i(s^t)$ is the wealth carried over from the previous period.

To prevent agents from rolling over debt forever (Ponzi schemes) in an infinite-horizon economy[2], we must impose a limit on how much they can borrow.

[1] In the macroeconomic literature, $a^i(s_{t+1}|s^t)$ and $a^i(s_{t+1}, s^t)$ are used interchangeably to denote the same state-contingent claim. The notation $a^i(s_{t+1}|s^t)$ emphasizes the conditional nature of the claim, mirroring the pricing kernel $Q(s_{t+1}|s^t)$. The notation $a^i(s_{t+1}, s^t)$ instead emphasizes the transition between nodes on the event tree. Because a future history is formally defined as $s^{t+1} = (s^t, s_{t+1})$, both expressions are frequently condensed to $a^i(s^{t+1})$ for brevity.

[2] Intuitively, in the absence of borrowing limits, an agent could borrow a large sum today to consume. Tomorrow, instead of using their own income to repay the debt, they could simply borrow an even larger amount to pay off the previous principal and interest. In an infinite-horizon economy, with no terminal date at which accounts must be settled, this process could in principle be repeated indefinitely, with the agent never sacrificing their own consumption. The *natural debt limit* is introduced precisely to rule out such

The *natural debt limit* imposes the loosest possible borrowing constraint such that the agent can repay their debt with certainty (in all possible future states).

Definition 1.6: Natural Debt Limit

The *natural debt limit*, namely the maximum allowable debt for agent i at node s^t, denoted by $A^i(s^t)$, is exactly the present discounted value of the agent's entire future endowment stream, evaluated using the sequential state prices:

$$A^i(s^t) = \sum_{\tau=t}^{\infty} \sum_{s^\tau | s^t} Q(s^\tau | s^t) y^i(s^\tau)$$

Thus, in order to rule out Ponzi schemes, we require that the agent's debt at node s^t cannot exceed this natural debt limit:

$$-a^i(s_{t+1}|s^t) \leq A^i(s_{t+1}), \quad \forall s_{t+1}.$$

Remark.

An "intuitive" worry might be, *"What if a streak of bad endowment realizations makes repayment impossible?"*. This is actually treating the natural debt limit $A^i(s^t)$ as a fixed, risk-free debt obligation (non-contingent debt). This intuition fails in complete markets because agents trade *state-contingent claims (Arrow securities)*, not fixed debt.

Mathematically, $A^i(s^t)$ does not correspond to an "expected" future income. It is simply the present market value of liquidating the agent's entire future endowment stream across all possible states at today's state prices $Q(s^\tau|s^t)$. When borrowing up to $A^i(s^t)$, the agent is not promising to repay a fixed amount regardless of tomorrow's state; rather, they are sacrificing their consumption and selling all of their state-specific endowments.

---

schemes, enforcing the requirement that any debt incurred must ultimately be backed by the agent's own future endowments.

Definition 1.7: Competitive Equilibrium (Sequential)

A competitive equilibrium consists of

- a set of pricing kernels $Q(s_{t+1}|s^t)$,
- a set of consumption allocations $\{c^i(s^t)\}$,
- a set of asset holdings $\{a^i(s^t)\}$, and
- given the natural debt limit $\{A^i(s^t)\}$

such that

- **Agent's Utility Max**:

  Given prices $Q(s_{t+1}|s^t)$ and the natural debt limit $\{A^i(s^t)\}$, the consumption allocation $\{c^i(s^t)\}_{t,s^t}$ and asset holdings $\{a^i(s^t)\}_{t,s^t}$ solve the following problem for each agent i:

$$\max_{\{c^i(s^t)\}_{t,s^t}} \sum_t \sum_{s^t} \beta^t \pi(s^t) u(c^i(s^t))$$
$$\text{s.t. } c^i(s^t) + \sum_{s_{t+1}|s^t} Q(s_{t+1}|s^t) a^i(s_{t+1}|s^t) \leq y^i(s^t) + a^i(s^t), \quad \forall t, \quad \forall s^t$$
$$a^i(s_{t+1}|s^t) \geq -A^i(s^{t+1}), \quad \forall s_{t+1}|s^t, \quad \forall t, \quad \forall s^t$$

- Market Clearing:
  - Goods Market: $\sum_i c^i(s^t) = \sum_i y^i(s^t)$ for all t and s^t.
  - Asset Market: $\sum_i a^i(s_{t+1}|s^t) = 0$ for all t and s^t.

The Lagrangian for the agent's problem is given by

$$\begin{aligned}
\mathcal{L} = &\sum_t \sum_{s^t} \beta^t \pi(s^t) u(c^i(s^t)) \\
&+ \sum_t \sum_{s^t} \lambda^i(s^t) \left(y^i(s^t) + a^i(s^t) - c^i(s^t) - \sum_{s_{t+1}|s^t} Q(s_{t+1}|s^t) a^i(s_{t+1}|s^t) \right) \\
&+ \sum_t \sum_{s^t} \sum_{s_{t+1}|s^t} \nu^i(s_{t+1}|s^t) \left(a^i(s_{t+1}|s^t) + A^i(s^{t+1}) \right),
\end{aligned}$$

where $\lambda^i(s^t)$ is the Lagrange multiplier for the sequential budget constraint at history s^t, and $\nu^i(s_{t+1}|s^t)$ is the Lagrange multiplier for the natural debt limit constraint at history s^t.

The FOC's are

- for $c^i(s^t)$:

$$\frac{\partial \mathcal{L}}{\partial c^i(s^t)} = \beta^t \pi(s^t) u'(c^i(s^t)) - \lambda^i(s^t) = 0 \quad \Longrightarrow \quad \lambda^i(s^t) = \beta^t \pi(s^t) u'(c^i(s^t)).$$

- for $a^i(s_{t+1}|s^t)$:

$$\begin{aligned}
&\frac{\partial \mathcal{L}}{\partial a^i(s_{t+1}|s^t)} = -\lambda^i(s^t)Q(s_{t+1}|s^t) + \lambda^i(s^{t+1}) + \nu^i(s_{t+1}|s^t) = 0 \\
\Longrightarrow \quad & \lambda^i(s^t)Q(s_{t+1}|s^t) = \lambda^i(s^{t+1}) + \nu^i(s_{t+1}|s^t).
\end{aligned}$$

Note that if the natural debt limit constraint is not binding, then $\nu^i(s_{t+1}|s^t) = 0$ and $a' > -A$. If the natural debt limit constraint is binding, then $\nu^i(s_{t+1}|s^t) > 0$ and $a' = -A$. However if $a' = -A$ for some period, then you cannot consume from this period forward. By the Inada condition, $\lim_{c\to 0} u'(c) = \infty$, $c = 0$ onwards is suboptimal. Hence, the natural debt limit constraint will never be binding in equilibrium, which implies that $\nu^i(s_{t+1}|s^t) = 0$ and $a' > -A$ for all t and s^t.

Having this result, we can rewrite the FOC as

$$Q(s_{t+1}|s^t) = \frac{\lambda^i(s^{t+1})}{\lambda^i(s^t)}.$$

And from the FOC for $c^i(s^t)$, we have had $\lambda^i(s^t) = \beta^t \pi(s^t) u'(c^i(s^t))$. Hence, we can rewrite the pricing kernel as

$$\begin{aligned}
Q(s_{t+1}|s^t) &= \frac{\beta^{t+1}}{\beta^t}\frac{\pi(s^{t+1})}{\pi(s^t)}\frac{u'(c^i(s^{t+1}))}{u'(c^i(s^t))} \\
&= \beta\pi(s^{t+1}|s^t)\frac{u'(c^i(s^{t+1}))}{u'(c^i(s^t))}.
\end{aligned}$$

Consequently,

$$Q(s_{t+1}|s^t) = \beta\pi(s^{t+1}|s^t)\frac{u'(c^i(s^{t+1}))}{u'(c^i(s^t))}.$$

In order to establish the equivalence between the sequential trading equilibrium and the time 0 trading equilibrium, we need to compare their FOC's. In the time 0 trading equilibrium, the FOC for $c^i(s^t)$ is given by

$$\beta^t \pi(s^t) u'(c^i(s^t)) = \mu^i q(s^t).$$

And this gives

$$\beta\pi(s^{t+1}|s^t)\frac{u'(c^i(s^{t+1}))}{u'(c^i(s^t))} = \frac{q(s^{t+1})}{q(s^t)} := q(s_{t+1}|s^t),$$

In order to make their FOC's coincide so that the two equilibria are equivalent, we need to set $Q(s_{t+1}|s^t) = q(s_{t+1}|s^t)$ for all t and s^t.[3] This implies that the pricing kernel in the sequential trading equilibrium is the same as the price of state-contingent claims in the time 0 trading equilibrium.

An additional requirement is that initial wealth assets must be zero in the sequential

[3] It is important to distinguish these two terms, as they originate from different market structures. $Q(s_{t+1}|s^t)$ is the *actual spot price* traded at node s^t in the sequential trading economy; it is the physical amount of s^t-goods an agent pays today for a one-period-ahead state-contingent claim. In contrast, $q(s_{t+1}|s^t) := q(s^{t+1})/q(s^t)$ is a relative price derived *algebraically* from the date-0 Arrow-Debreu economy, in which no actual trading occurs at date t. Setting $Q = q$ ensures that the sequential spot markets exactly replicate the relative valuations embedded in the date-0 prices—this is the mathematical content of the equivalence between the two market structures.

trading equilibrium, which means that $a^i(s_0) = 0$ for all i. This is because in the time 0 trading equilibrium, all trades occur at time 0, and thus there are no initial assets or debts. If there are initial assets or debts in the sequential trading equilibrium, then the two equilibria cannot be equivalent since the initial conditions are different.

With the above two requirements, we can establish the equivalence between the sequential trading equilibrium and the time 0 trading equilibrium.

Remark (Asset Pricing).

Assume $\{d(s^t)\}_t$ is a stream of claims of consumption. The price of the consumption stream at time 0 is given by

$$p(s_0) = \sum_t \sum_{s^t} q(s^t|s_0)d(s^t).$$

After a realization of history s^τ, the price of the consumption stream is given by

$$p(s^\tau) = \sum_{t\geq\tau} \sum_{s^t} q(s^t|s^\tau)d(s^t),$$

where $q(s^t|s^\tau) = \frac{q(s^t|s_0)}{q(s^\tau|s_0)}$, and by the previous section, we have

$$q(s^t|s^\tau) = \frac{q(s^t|s_0)}{q(s^\tau|s_0)} = \beta^{t-\tau}\pi(s^t|s^\tau)\frac{u'(c^i(s^t))}{u'(c^i(s^\tau))}.$$

Of the same logic, the price of a one-period return is given by

$$p(s^t) = \sum_{s_{t+1}|s^t} q(s_{t+1}|s^t)d(s_{t+1}) = \sum_{s_{t+1}|s^t} \beta\pi(s_{t+1}|s^t)\frac{u'(c^i(s_{t+1}))}{u'(c^i(s^t))}d(s_{t+1}).$$

Define the *gross return of the asset* as

$$R_{t+1} = \frac{d(s_{t+1})}{p(s^t)}.$$

Define the *stochastic discount factor* as[a]

$$m_{t+1} = \beta\frac{u'(c^i(s_{t+1}))}{u'(c^i(s^t))}.$$

Then we can rewrite the price of the one-period return as

$$1 = \sum_{s_{t+1}|s^t} \pi(s_{t+1}|s^t)m_{t+1}R_{t+1} = \mathbb{E}_t[m_{t+1}R_{t+1}|s^t].$$

[a]The stochastic discount factor represents the agent's subjective valuation of one unit of consumption tomorrow in state s_{t+1} relative to one unit today. It is essentially the Intertemporal Marginal Rate of Substitution (IMRS).

Suppose a bad economic shock occurs tomorrow. Consumption c_{t+1} will be low, meaning marginal utility $u'(c_{t+1})$ will be extremely high (due to diminishing marginal utility). Thus, the SDF m_{t+1} takes a very large value in bad states. This implies that an asset paying off well during bad times is heavily weighted by the SDF, commanding a higher price today. This is the fundamental source of risk premiums in modern asset pricing.

1.4 Recursive Competitive Equilibrium

In the sequential trading setup, agents choose sequences of consumption and asset holdings across a constantly expanding event tree of histories s^t. As $t \to \infty$, tracking the entire history makes the problem infinitely dimensional. The goal of a *recursive competitive equilibrium* is to collapse this infinite-dimensional sequence problem into a time-invariant functional equation (the Bellman equation). To do this, we must ensure that the past only influences the future through a finite set of current "state variables" (e.g., current wealth a and current state s). This requires strict stationarity restrictions on the fundamental economic environment.

To guarantee that the problem can be written recursively, we impose the following assumptions:

Assumption 1.8: Markovian and Time-Invariant Fundamentals

- **Markovian Transitions:** Endowments are governed by a first-order Markov process. The probability of tomorrow's state depends *only* on today's state, making the history prior to t irrelevant for forecasting the future:

$$\Pr(s_{t+1} = s' | s_t = s) = \pi(s'|s), \quad \forall s, \quad \forall s'.$$

- **Time-Invariant Endowments:**[a] Household endowments depend only on the current state s, not on the date t or the path taken to get there:

$$y^i(s^t) = y^i(s_t), \quad \forall i, \quad \forall t.$$

[a]This means the current physical state s is a *sufficient statistic* for the endowment. You only need to know the current state, instead of the whole history s^t or the time t, to determine current income.

Because the fundamental environment (endowments and transition probabilities) is now memoryless and stationary, the equilibrium consumption allocations, asset holdings, and prices will also depend only on the current state. Thus, we can safely drop all t subscripts from the sequential budget constraint. The constraint simplifies to:

$$c(s) + \sum_{s'} Q(s'|s)a(s') \leq y(s) + a(s), \quad \forall s.$$

We can now formally write the household's dynamic programming problem. The state variables are individual wealth a and the current exogenous state s. The Bellman equation is given by:

$$\begin{aligned} V^i(a,s) &= \max_{c(s),\hat{a}(s')} \left\{ u^i(c(s)) + \beta \sum_{s'} \pi(s'|s) V^i(\hat{a}(s'), s') \right\} \\ \text{s.t. } & c(s) + \sum_{s'} Q(s'|s)\hat{a}(s') \leq y(s) + a, \quad \forall s \\ & \hat{a}(s') \geq -A^i(s'), \quad \forall s' \\ & c(s) \geq 0, \quad \forall s. \end{aligned}$$

where $V^i(a,s)$ is the value function of the household, representing the maximum expected discounted utility that the household with current assets a and current state s can achieve.

From the Bellman equation, we can derive the policy functions for consumption and asset holdings, denoted by $c = h^i(a,s)$ and $\hat{a}(s') = g^i(a,s,s')$, respectively.

Definition 1.9: Recursive Competitive Equilibrium

A recursive competitive equilibrium consists of

- pricing kernels: $Q(s'|s)$,
- policy functions: $\{h^i(a,s), g^i(a,s,s')\}_i$,
- value functions: $\{V^i(a,s)\}_i$,
- initial distribution of wealth: $\{a_0^i\}_i$, and
- given the natural debt limits $\{A^i(s)\}_i$

such that

- Given the pricing kernel $Q(s'|s)$, the natural debt limits $A^i(s)$ and the initial distribution of wealth a_0^i, the policy functions $h^i(a,s)$ and $g^i(a,s,s')$ solve the household's Bellman equation.
- Market Clearing:
  - Goods Market: $\sum_i h^i(a^i,s) = \sum_i y^i(s)$ for all s.
  - Asset Market: $\sum_i g^i(a^i,s,s') = 0$ for all s and s'.
- State-by-state borrowing constraint:

$$A^i(s) = y^i(s) + \sum_{s'} Q(s'|s) A^i(s'), \quad \forall i, \quad \forall s.$$

Remark (The Necessity of the Recursive Debt Limit Constraint).

- In the sequential problem, the natural debt limit is defined as an infinite sum of discounted future endowments. In a recursive framework, we drop the time dimension and cannot explicitly sum to infinity. This equation is the exact *recursive formulation* of that infinite sum. It acts as a functional equation that endogenously solves for the unknown function $A^i(s)$.

- The debt limit $A^i(s)$ is not an exogenous parameter; it strictly depends on the endogenous equilibrium pricing kernel $Q(s'|s)$. If we omit this condition, the household's constraint $\hat{a}(s') \geq -A^i(s')$ is left undefined, and the model cannot uniquely rule out Ponzi schemes.

The FOC for the Bellman equation is given by

- for c:
$$u'(c) = \lambda(s),$$
where $\lambda(s)$ is the Lagrange multiplier for the budget constraint at state s.

- for $\hat{a}(s')$:
$$\beta\pi(s'|s)V_a(\hat{a}, s') - \lambda(s)Q(s'|s) + \nu(s, s') = 0,$$
where $\nu(s, s')$ is the Lagrange multiplier for the natural debt limit constraint at state s'.

Note that $\nu > 0$ if the borrowing constraint is binding. Then at any future state s', the agent can only have zero consumption, $c(s') = 0$. By the Inada condition, $c(s') = 0$ is suboptimal. Hence, the natural debt limit constraint will never be binding in equilibrium, which implies that $\nu(s, s') = 0$ for all s and s'.

The envelope condition for the Bellman equation is given by

$$V_a(a, s) = u'(c).$$

Combining the conditions above, we have

$$\beta\pi(s'|s)u'(c(s')) - u'(c(s))Q(s'|s) = 0 \implies Q(s'|s) = \beta\pi(s'|s)\frac{u'(c(s'))}{u'(c(s))},$$

which is the same as the pricing kernel in the sequential trading equilibrium.[4] This implies that the recursive competitive equilibrium and the sequential trading equilibrium are equivalent under our assumptions.

1.5 Recursive Social Planner's Problem

*Why do we need the **Recursive Social Planner's Problem** and the concept of **promised utility** (v) as a state variable, especially when the Sequential Planner already easily solves for the optimal allocation?*

This transition represents a fundamental paradigm shift in how we approach dynamic optimization. In the previous social planner's problem, the planner stands at time 0, assigns a fixed Pareto weight, and dictates an infinite sequence of allocations for all future contingencies. The "contract" is written once and never revisited.

The recursive planner, however, shifts the perspective from "static allocation over infinite time" to "dynamic utility delivery." By using promised utility v as the state variable, the

[4] Recall that in the sequential trading, $Q(s_{t+1}|s^t) = \beta\pi(s_{t+1}|s^t)\frac{u'(c^i(s^{t+1}))}{u'(c^i(s^t))}$. Applying our assumptions, the conditional probability collapses to $\pi(s_{t+1}|s_t)$, and the consumption allocations collapse to $c^i(s_{t+1})$ and $c^i(s_t)$. Thus, the price simplifies to a time-invariant function mapping today's state s to tomorrow's state s': $Q(s'|s) = \beta\pi(s'|s)\frac{u'(c^i(s'))}{u'(c^i(s))}, \forall s, \forall s'$.

planner collapses the entire infinite future into a single sufficient statistic: the "utility debt" currently owed to the agent. The planner's problem is no longer about choosing an entire physical consumption path. Instead, it is a sequential constrained optimization problem: ***how to deliver this promised utility v as efficiently as possible.***

Each period, the planner wakes up, looks at the ledger (v), and faces a one-step tradeoff: how to split this obligation into today's physical consumption (c_s) and tomorrow's renewed promise (w_s). This approach changes the mathematical object we are operating on: we are no longer choosing sequences of goods, but recursively trading off utility across time. In this frictionless baseline the planner optimally keeps the promise constant ($w_s = v$, mirroring a fixed Pareto weight), but the "ledger" perspective gives us a much deeper understanding of the mechanics of time and commitment that becomes essential once we add frictions in subsequent chapters.

Assume:

- Two agents with constant aggregate endowment:

$$y_t^1 = s_t, \quad y_t^2 = 1 - s_t, \quad \forall t.$$

- The endowment process is i.i.d.:

$$\pi(s^t) = \pi(s_0)\pi(s_1)\cdots\pi(s_t) = \prod_{\tau=0}^{t} \pi(s_\tau), \quad \forall t.$$

- s_t has a discrete distribution: $s_t \in [\bar{s}_1, \cdots, \bar{s}_S]$.

We assume *ex-ante* timing, meaning the decision at time t is made before the realization of s_t. The social planner delivers a promised utility stream v for Agent 1 at time t such that

$$\sum_s \pi(s)\left(u(c_s) + \beta w_s\right) = v,$$

where c_s is the consumption at state s and w_s is the continuation value for Agent 1 at state s, i.e., if the current state is s, then the social planner will deliver a promised utility stream w_s for Agent 1 from the next period onward.

Before formally setting up the recursive problem, we must define the planner's value function, $P(v)$. Let $P(v)$ represent the *Pareto frontier*: the maximum expected discounted utility the social planner can deliver to Agent 2, given that Agent 1 is guaranteed a promised utility of at least v. Since the aggregate endowment is constant at 1, if Agent 1 consumes c_s, Agent 2 must consume the residual $1 - c_s$. The planner's objective is to maximize Agent 2's utility subject to fulfilling the utility promise v to Agent 1. So the recursive Pareto problem is given by

$$\begin{aligned} P(v) = \max_{\{c_s, w_s\}_s} &\sum_s \pi(s)\left(u(1 - c_s) + \beta P(w_s)\right) \\ \text{s.t. } &\sum_s \pi(s)\left(u(c_s) + \beta w_s\right) \geq v. \end{aligned}$$

The Lagrange for the problem is given by

$$\mathcal{L} = \sum_s \pi(s)\left(u(1-c_s) + \beta P(w_s)\right) + \theta\left(\sum_s \pi(s)\left(u(c_s) + \beta w_s\right) - v\right),$$

where θ is the Lagrange multiplier for the promised utility constraint.

The FOC's are given by

- for c_s:

$$-u'(1-c_s) + \theta u'(c_s) = 0 \quad \Longrightarrow \quad \theta = \frac{u'(1-c_s)}{u'(c_s)}.$$

- for w_s:

$$\beta P'(w_s) + \theta\beta = 0 \quad \Longrightarrow \quad P'(w_s) + \theta = 0.$$

And the envelope condition is given by

$$P'(v) = -\theta.$$

This provides a beautiful economic interpretation for the Lagrange multiplier θ. Since the Pareto frontier $P(v)$ represents Agent 2's maximum utility given Agent 1's promised utility v, its derivative $P'(v)$ must be strictly negative. It measures the marginal cost to Agent 2 of delivering one additional unit of utility to Agent 1. Thus, $-P'(v) = \theta$ is exactly the *marginal rate of utility substitution* between the two agents.

From those conditions, we have

$$P'(w_s) = P'(v) \quad \Longrightarrow \quad w_s = v, \quad \forall s.$$

And we also have

$$\frac{u'(1-c_s)}{u'(c_s)} = \theta = -P'(v).$$

This implies that the consumption is constant over time and across states: $c_s = c$ for all s, which indicates complete risk-sharing of the idiosyncratic risk. Notably, this recursive equilibrium corresponds to some allocation in the sequential social planner's problem.

Intuitively, these results say that the planner keeps this "utility exchange rate" constant across all future states s. If the exchange rate of utility never fluctuates, the physical allocation of consumption must also remain perfectly smooth, which is the perfect-risk-sharing conclusion $c_s = c$.

Remark (Chapter Summary).

- **Three equivalent representations of the same equilibrium.** The complete-markets allocation can be characterized as (i) the social planner's optimum with Pareto weights λ^i, (ii) a date-0 Arrow–Debreu trading equilibrium with state-contingent prices $q(s^t)$, and (iii) a sequential-trading equilibrium with one-period-ahead Arrow securities priced by the kernel $Q(s'|s)$. The First and Second Welfare Theorems supply the bridges between them.

- **The natural debt limit $A^i(s^t)$.** Borrowing up to the present value of future endowments is the loosest possible constraint compatible with feasibility. Under Inada conditions it never binds in equilibrium, so it can be written down without affecting the analysis.

- **Recursive formulation needs stationarity.** To collapse the infinite-history problem into a Bellman equation in (a, s), we impose that endowments and transitions are Markov and time-invariant.
- **Recursive social planner via promised utility.** Using v (or w_s) as the state variable lets the planner deliver utility "on a ledger." In the frictionless setting the optimal policy keeps v constant, recovering perfect risk sharing $c_s = c$.
- **The benchmark to keep in mind.** With complete markets and no aggregate risk, idiosyncratic shocks are insured away completely. Subsequent chapters break this benchmark in two distinct ways—exogenously (Chapter 2) and endogenously through participation frictions (Chapter 3).

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 Huggett (1993) 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

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

Here ϕ is the *ad-hoc* borrowing constraint. Different from the natural borrowing limit, it might be the case that $\phi \le \frac{\min\{y_1, \cdots, y_N\}}{R} = \frac{y_1}{R}$ (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

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

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

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

$$\Longrightarrow 0 = 0 + \frac{\psi(1-\psi)}{R-\psi}(\bar{y}-\bar{y}) + \frac{\gamma(\beta-1)}{(R-\psi)^2}\frac{\sigma^2}{2} + \frac{\ln \beta R}{\gamma(R-1)}$$

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.

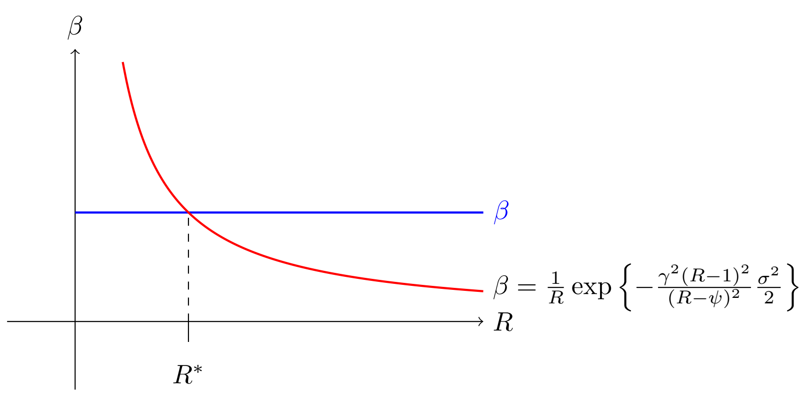

From the graphical intuition, we can see that the general equilibrium interest rate R^* increases when there is

- 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.,

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

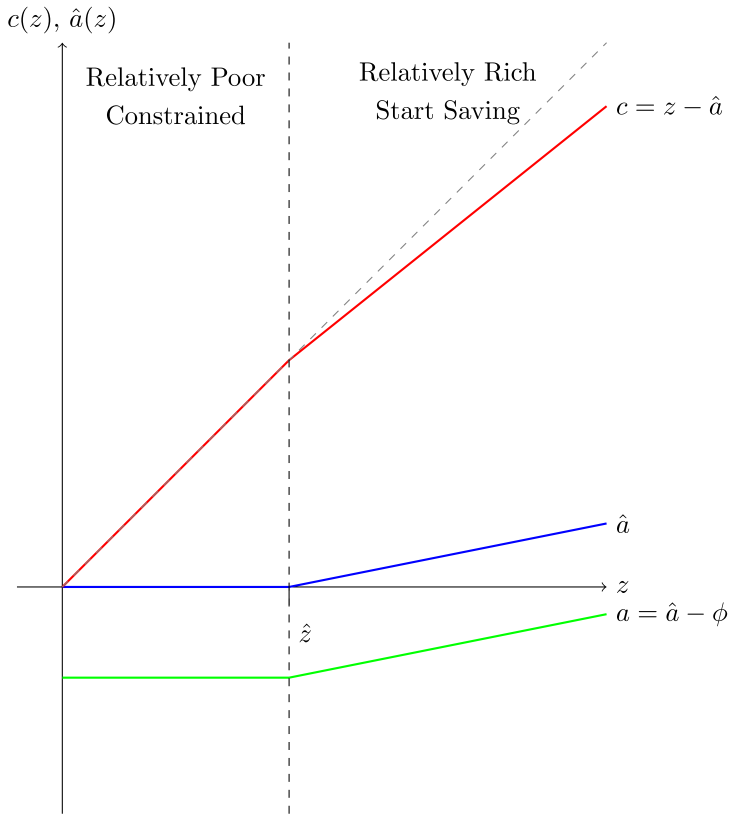

Thus, this allows us to have a general picture of the ergodic distribution of assets in the economy.[10]

[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.*

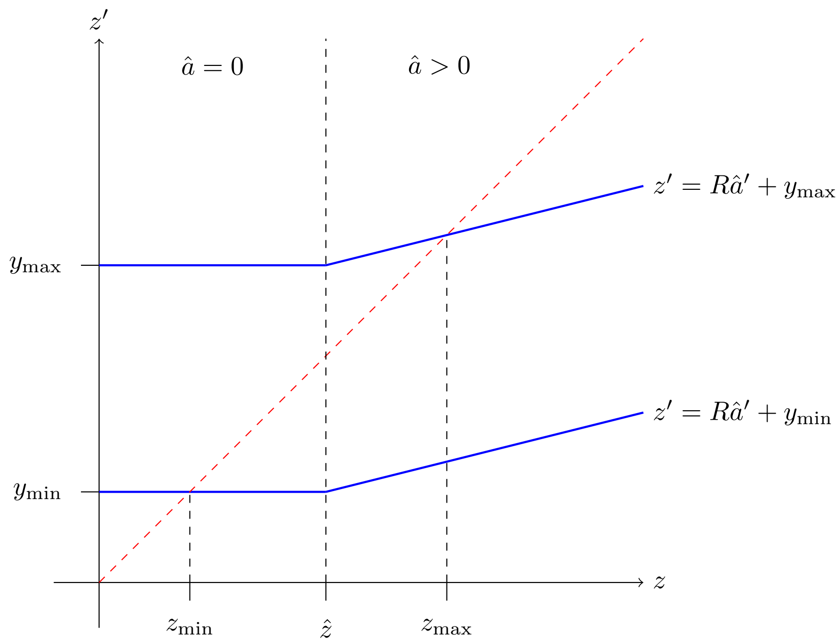

2.3.4 Compute the Stationary Distribution of Assets

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:

$$r_k = F_K(K,L),$$
$$w = F_L(K,L).$$

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

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

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).

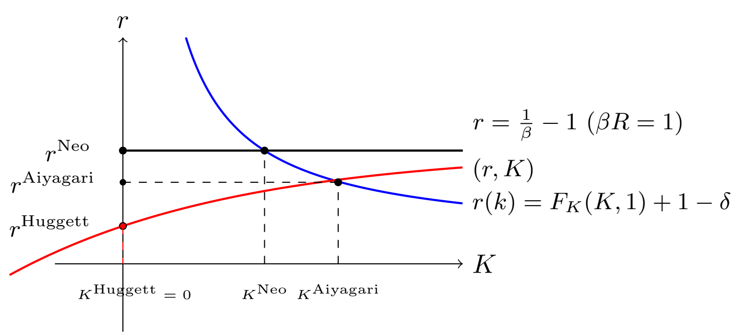

- **The Blue Curve (Capital Demand):** This represents the firm's first-order condition, $r = F_K(K,1) - \delta$. It is downward-sloping due to the diminishing marginal product of capital.

- **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 Krusell and Smith (1998) 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

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

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).

Chapter 3

Endogenously Incomplete Markets

Remark (Notation in This Chapter).

Symbol	Meaning
V^{aut}	Autarky continuation value (consume own endowment forever)
v, w_s	Promised lifetime utility today; continuation promise after state s
$P(v)$	Money lender's profit frontier in the one-sided LoC model
Δ, $\bar{\Delta}$	Surplus utility above autarky; upper bound of Δ
$Q(\Delta, s)$	Pareto frontier in the two-sided LoC model
$\bar{c}, \underline{c}$	Upper / lower bound of the consumption interval at a given state
$\lambda_s, \theta_{s'}$	Lagrange multipliers on the participation constraints
μ	Multiplier on the promise-keeping constraint
$W(s^t)$	Agent's wealth $\equiv$ PV of future endowments (Bulow–Rogoff)
$D(s^t)$	Agent's debt $\equiv$ PV of future payments (Bulow–Rogoff)
k	Debt-to-wealth limit ratio (Bulow–Rogoff)
$a(s^t), g(s^{t+1})$	Cash-in-advance saving and state-contingent payoff (Bulow–Rogoff)

3.1 One-Sided Lack of Commitment

Assume:

- Money lender (ML):
  - risk neutral.
  - can borrow and lend money at the interest rate R such that $\beta R = 1$.
  - contracts are state-dependent.
  - all borrowings and savings can only go through the money lender.
- Consumer:

– preferences:

$$\mathbb{E}\left[\sum_{t=0}^{\infty}\beta^t u(c_t)\right],$$

where u is assumed to be strictly increasing, strictly concave, and twice continuously differentiable.

– stochastic endowment: $y(s^t)$ that is i.i.d. with each state s^t occurring with probability $\pi(s^t)$. And the space of $y(s^t)$ is finite such that

$$y_1 < y_2 < \cdots < y_N.$$

– Consumers cannot commit to a contract: if the consumer walks out of the contract at some point, then the consumer will not be able to borrow or lend from the money lender in the future. That is, from the period onwards, the consumer can only consume their endowment $y(s^t)$ in each period ("*Autarky*").

Before solving the complex model with lack of commitment, it is standard practice to establish the frictionless "First-Best" (efficient) benchmark.

Claim: First-Best Benchmark

The efficient allocation is such that the consumer consumes the average endowment at each period and each state, i.e.,

$$\sum_t \beta^t u\left(\sum_s \pi_s y_s\right) = \frac{1}{1-\beta} u\left(\sum_s \pi_s y_s\right).$$

Thus, the profit of the money lender is given by

$$\mathbb{E}\left[\sum_{t=0}^{\infty}\beta^t\left(y(s^t)-c_t\right)\right] = \mathbb{E}\left[\sum_{t=0}^{\infty}\beta^t\left(y(s^t)-\sum_s \pi_s y_s\right)\right].$$

Why is it efficient? The economy consists of a risk-neutral Money Lender (ML) and a strictly risk-averse consumer. A fundamental result in microeconomics is that Pareto efficiency requires the risk-neutral party to fully absorb all risk. Therefore, the ML should provide *perfect insurance*, keeping the consumer's consumption perfectly smooth across all states and time. To ensure the contract is feasible (and the ML breaks even in expectation), this constant consumption level must exactly equal the expected (average) endowment: $c_t = \sum_s \pi_s y_s := \bar{y}$.

But later we will see that this efficient allocation is not sustainable under the lack of commitment friction. Intuitively, this perfect insurance contract requires the consumer to make net payments to the ML in good states ($y_s > \bar{y}$) and receive transfers in bad states ($y_s < \bar{y}$). However, under *lack of commitment*, when a good state realizes (e.g., the best state $y_N > \bar{y}$), the consumer's autarky payoff $u(y_N)$ at the current period strictly dominates the contract's smoothed payoff $u(\bar{y})$, which makes the consumer want to walk out of the contract and consume the endowment y_N instead.

3.1.1 Self-Enforcing Contract

Due to the lack of commitment, the ML will find a contract that is *self-enforcing* (also called *sustainable*), which means that the consumer will not want to walk out of the contract at any point in time.

Assume the ML is offering a stream of consumption $\{c(s^t)\}_{t=0}^{\infty}$ to the consumer.

Under the contract, the consumer's utility is given by

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

For the money lender, the profit is given by

$$P = \sum_t \sum_{s^t} \beta^t \pi(s^t) \left(y(s^t) - c(s^t) \right).$$

If at some point the consumer decides to walk out of the contract, then the consumer's utility will be given by

$$\begin{aligned} & u(y(s)) + \beta \sum_{s'} \pi(s'|s) u(y(s')) + \beta^2 \sum_{s''} \pi(s''|s') u(y(s'')) + \cdots \\ =& u(y(s)) + \frac{\beta}{1-\beta} \sum_{s'} \pi(s'|s) u(y(s')) \\ :=& u(y(s)) + \beta V^{\text{aut}}. \end{aligned}$$

For a contract to be self-enforcing, the ML must offer at least the utility that the consumer can get by walking out of the contract, for all t and s^t.

The natural question is then, what is the optimal contract?

The ML wants to maximize the profit P by choosing the contract $\{c(s^t)\}_{t=0}^{\infty}$ subject to the self-enforcing constraint.

To analyze this question, we make an important assumption about the timing. Specifically, we assume *ex-ante* timing, which means that at each state, the consumer wakes up and makes decisions based on the expected realizations, and then the state (and thus the endowment realization) is revealed.

The ML's problem is given by

$$\begin{aligned} P(v) = \max_{c_s, w_s} & \sum_s \pi_s \left[y(s) - c(s) + \beta P(w) \right] \\ \text{s.t. } & \sum_s \pi_s \left[u(c_s) + \beta w_s \right] \geq v, \\ & u(c_s) + \beta w_s \geq u(y_s) + \beta V^{\text{aut}}, \quad \forall s, \end{aligned}$$

where the first constraint is the *promise-keeping constraint*, and the second constraint is the *participation constraint* (or, the self-enforcing constraint from the perspective of the ML).

Here we in addition assume that

$$\begin{aligned} w_s &\in \left[V^{\text{aut}}, V^{\max} \right], \quad \forall s, \\ c_s &\in \left[c_{\min}, c_{\max} \right], \quad \forall s. \end{aligned}$$

And we give the following claim without proof:

Claim

- $P(v)$ is strictly decreasing in v.
- $P(v)$ is strictly concave in v.
- $P(v)$ is continuously differentiable in v.

The Lagrange is given by

$$\begin{aligned}\mathcal{L} = &\sum_s \pi_s \left[y_s - c_s + \beta P(w_s)\right] \\ &+ \mu \left[\sum_s \pi_s \left[u(c_s) + \beta w_s\right] - v\right] \\ &+ \sum_s \lambda_s \left[u(c_s) + \beta w_s - u(y_s) - \beta V^{\text{aut}}\right].\end{aligned}$$

The FOC is given by

- w.r.t. c_s:
$$-\pi_s + \mu \pi_s u'(c_s) + \lambda_s u'(c_s) = 0.$$
- w.r.t. w_s:
$$\beta \pi_s P'(w_s) + \mu \beta \pi_s + \lambda_s \beta = 0.$$

The envelope condition is given by

$$P'(v) = -\mu.$$

From the three conditions above, we can obtain the following results:

- $u'(c_s) = -1/P'(w_s)$.
- $P'(v) = \lambda_s/\pi_s + P'(w_s)$.
- $1/u'(c_s) = \lambda_s/\pi_s + 1/u'(c_{s-1})$.

We analyze the problem by considering two cases:

- Case 1: Participation constraint is not binding.

  Then $\lambda_s = 0$. This immediately implies
$$P'(v) = P'(w_s), \quad 1/u'(c_s) = 1/u'(c_{s-1}).$$
  Since $P'(\cdot)$ and $u'(\cdot)$ are strictly increasing, we have $v = w_s$ and $c_s = c_{s-1}$. So when the participation constraint does not bind, the consumption and promised value are just constants and do not depend on the state s.
- Case 2: Participation constraint is binding for some s.

Then $\lambda_s > 0$. This implies

$$\begin{cases} P'(v) > P'(w_s) \implies v < w_s, \\ 1/u'(c_s) > 1/u'(c_{s-1}) \implies c_{s-1} < c_s. \end{cases}$$

Intuition: The consumer has incentives to walk out of the contract when they are in a good state (i.e., when $y(s)$ is high). To prevent the consumer from walking out of the contract, the ML needs to offer a higher promised value w_s and a higher consumption c_s in the good state than in the bad state.

In conclusion, the consumption is always non-decreasing: strictly increases when the participation constraint binds, and remains constant when the participation constraint does not bind. When the participation constraint binds (i.e., the consumer wakes up in a good state), the ML incentivizes the consumer to stay in the contract by offering a higher promised value and a higher consumption from then on (although the consumer will then need to forgo some of their current endowment since $c_s < y_s$[1]).

Motivated by the analysis above, for a given promised value v, there exists $\bar{y}(v)$ such that the participation constraint binds when $y(s) > \bar{y}(v)$ and does not bind when $y(s) < \bar{y}(v)$. In order to find such $\bar{y}(v)$, we can solve the following equations:

$$\begin{cases} w_s = v \\ c_s = c_{s-1} \\ u(c_s) + \beta w_s = u(y(s)) + \beta V^{\text{aut}}. \end{cases}$$

Below we give graphical illustrations of the above analysis. When $\lambda_s = 0$, $c_s = c_{s-1}$ and $w_s = v$. When $\lambda_s > 0$, $c_s := \tilde{c}(s)$ such that $\tilde{c}(s) > c_{s-1}$ and $w_s := \tilde{w}(s)$ such that $\tilde{w}(s) > v$.

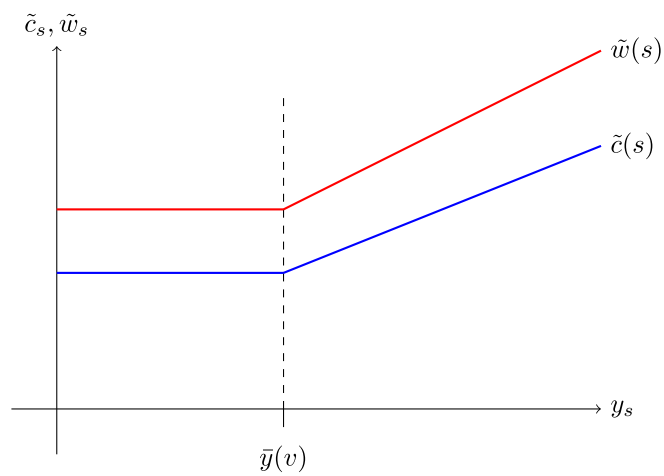

Consider the following consumption dynamics. We assume that the participation constraint does not bind at y_1, but binds at $y_2, \cdots, y_N$.

[1]Here we show that $c_s < y_s$ for all state s. In order to make the consumer participate in the contract, the ML must make the promised value $v \geq V^{\text{aut}}$. When the participation constraint binds, we have $u(c_s) + \beta w_s = u(y(s)) + \beta V^{\text{aut}}$. Since $w_s > v \geq V^{\text{aut}}$, we have $u(c_s) < u(y(s))$, which implies $c_s < y_s$. This happens when the participation constraint binds, so when the participation constraint does not bind, we still have $c_s < y_s$. Intuitively, this is true because the ML needs to make a positive profit from the contract, so the consumption stream offered to the consumer must be less than their endowment in each state.

t	1	2	3	4	5	6	7
y_t	y_1	y_1	y_2	y_1	y_N	y_2	y_1

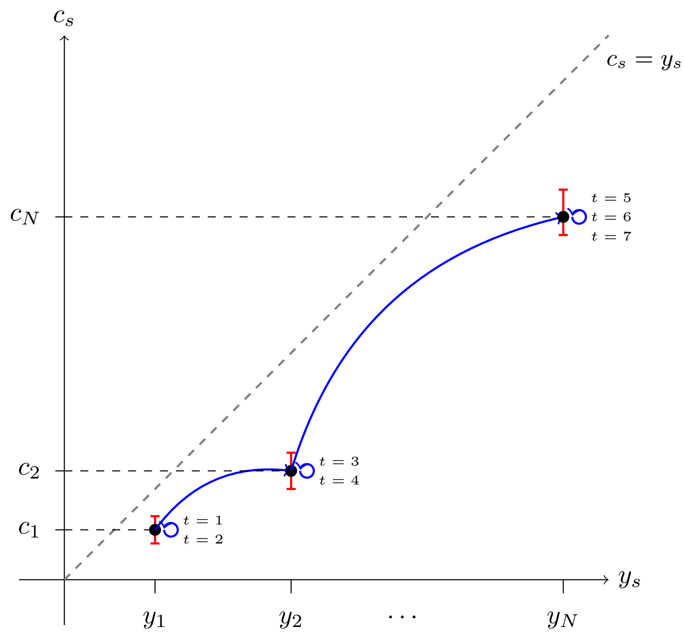

Initially, the consumer wakes up in a bad state with endowment y_1 and the participation constraint does not bind, so the consumer will consume c_1. At $t = 2$, the consumer wakes up in the same state with endowment y_1, so the consumption remains unchanged. At $t = 3$, the consumer wakes up in a good state with endowment y_2, and the participation constraint binds. In order to incentivize the consumer to stay in the contract, the ML needs to offer a higher promised value and a higher consumption c_2 such that $c_2 > c_1$. At $t = 4$, the consumer wakes up in the same state with endowment $y_1 < y_2$, so the participation constraint does not bind, and then the consumption remains unchanged at c_2. At $t = 5$, the consumer wakes up in a better state with endowment y_N, and the participation constraint binds again. The ML then offers a higher promised value and a higher consumption c_N such that $c_N > c_2$ to keep the consumer in the contract. At $t = 6$ and 7, although the consumer wakes up in relatively bad states, the consumption remains unchanged (in particular, does not decrease) at c_N since the participation constraint does not bind.

3.1.2 Profit of the ML

Ex-ante, the ML would offer

$$v = V^{\text{aut}}.$$

If y_s happens to be y_1, and $v = V^{\text{aut}}$, we consider the case when the participation constraint binds:

$$u(c_1) + \beta v = u(y_1) + \beta V^{\text{aut}}.$$

This implies the ML must offer the consumer $c_1 = y_1$ to make them stay in the contract. Then the ML will make zero profit from the consumer when $y_s = y_1$. (From the reasoning we also know that $\bar{y}(V^{\text{aut}}) = y_1$.)

When y_s happens to be y_N at some point, the consumption will be c_N from then on. In other words, the consumption will converge to $c_N := \bar{c}$ once the best state y_N is realized.

Again, the ML will offer c_N to make the consumer just stay in the contract when $y_s = y_N$:

$$u(\bar{c}) + \beta \frac{u(\bar{c})}{1-\beta} = u(y_N) + \beta V^{\text{aut}} = u(y_N) + \frac{\beta}{1-\beta} \sum_s \pi_s u(y_s).$$

Previously we have shown that $c_s < y_s$ for all state s, so we have $\bar{c} < y_N$. This implies that

$$u(\bar{c}) > \sum_s \pi_s u(y_s).$$

Namely, the value of the contract is higher than the (expected) value of autarky.

Once the best state y_N is reached, the profit of the ML is given by

$$\begin{aligned} P &= y_N - \bar{c} + \frac{\beta}{1-\beta} \sum_s \pi_s (y_s - \bar{c}) \\ &= y_N + \frac{\beta}{1-\beta} \sum_s \pi_s y_s - \frac{\bar{c}}{1-\beta}. \end{aligned}$$

Claim

The profit of the ML is positive after the best state y_N is reached.

Proof for Claim.

The profit of the ML is positive if and only if

$$y_N + \frac{\beta}{1-\beta} \sum_s \pi_s y_s > \frac{\bar{c}}{1-\beta}.$$

We prove this by contradiction. Suppose not, then we have

$$\begin{aligned} & y_N + \frac{\beta}{1-\beta} \sum_s \pi_s y_s \le \frac{\bar{c}}{1-\beta} \\ \iff & (1-\beta) y_N + \beta \sum_s \pi_s y_s \le \bar{c} \\ \iff & u\left((1-\beta) y_N + \beta \sum_s \pi_s y_s \right) \le u(\bar{c}) = (1-\beta) u(y_N) + \beta \sum_s \pi_s u(y_s), \\ \iff & u\left((1-\beta) y_N + \beta \sum_s \pi_s y_s \right) \le (1-\beta) u(y_N) + \beta \sum_s \pi_s u(y_s). \end{aligned}$$

where the equality in the third line holds from the participation constraint:

$$u(\bar{c}) + \beta \frac{u(\bar{c})}{1-\beta} = u(y_N) + \frac{\beta}{1-\beta} \sum_s \pi_s u(y_s) \implies u(\bar{c}) = (1-\beta) u(y_N) + \beta \sum_s \pi_s u(y_s).$$

But the last inequality holds if and only if u is convex, by Jensen's inequality. This contradicts the assumption that u is strictly concave. Therefore, we must have $P > 0$

after the best state y_N is reached.

Remark (Intuition).

- Intuitively, the ML is making positive profit because of risk sharing. The consumer is risk averse, while the ML is risk neutral. After hitting the best state, the consumer is insured against the idiosyncratic risk through the contract, so the consumer is willing to pay a positive amount (by sacrificing the current consumption of $y_N - c_N > 0$) to the ML to get insurance. The ML then makes positive profit from the contract (*risk premium*).
- The agent owes money to the ML in y_N state, but receives net transfers (insurance) from the ML in future bad states. The ML makes positive profit from the contract because the consumer is risk averse and thus values the insurance provided by the contract.

Remark (Why focus only on y_1 and y_N? What about the intermediate states?).

In the analysis of the one-sided lack of commitment model, the notes explicitly target y_1 and y_N because they serve as the two critical **boundary conditions** of the dynamic contract. The intermediate states are omitted not because they do not exist, but because characterizing the boundaries is sufficient to demonstrate the economic mechanism of the contract.

In the intermediate states, the single-period profit $y_s - c_s$ is transitional and highly history-dependent. The ML might suffer temporary losses (paying out insurance when the agent draws bad shocks) or earn small profits. To establish that the contract is self-enforcing and profitable *ex-ante*, however, it suffices to show that the ML breaks even at the bottom of the state space (y_1) and strictly collects a risk premium once the top (y_N) is reached. The intermediate states then take care of themselves.

3.2 Bulow-Rogoff Contracts

This model deviates from the standard one-sided lack of commitment model. In standard models (e.g., Eaton-Gersovitz 1981), the punishment for default is a permanent exclusion from future borrowing, which acts as a sufficient threat to sustain debt. This model introduces a critical loophole: it allows the agent to ***save*** *in complete financial markets after defaulting. We will prove that if an agent has the ability to save after default, borrowing through the Money Lender (ML) becomes completely infeasible (i.e., the credit limit inevitably collapses to zero).*

We assume

- **Agent:** Risk-averse with strictly increasing utility $u(\cdot)$. Receives a stochastic endowment stream $y(s)$.
- **Money Lender (ML):** Risk-neutral and competitive. Has access to complete financial markets and discounts the future at a risk-free gross interest rate $R > 1$.

Definition 3.1: Wealth

The agent's wealth at node s^t, denoted by $W(s^t)$, is defined as the expected present discounted value of all future endowments:

$$W(s^t) = \sum_{\tau \geq t} \sum_{s^\tau | s^t} \pi(s^\tau | s^t) \frac{y(s^\tau)}{R^{\tau - t}}$$

In addition, we assume $W(s^t)$ is finite (which requires R to be sufficiently high).

Definition 3.2: Contract

A contract is defined by a stream of payments $p(s^t)$ from the agent to the ML.

- $p(s^t) > 0$: The agent is paying back the ML.
- $p(s^t) < 0$: The agent is borrowing (receiving inflows from the ML).

Definition 3.3: Debt

The agent's debt at node s^t, denoted by $D(s^t)$, is defined as the expected present discounted value of all future payments:

$$D(s^t) = \sum_{\tau \geq t} \sum_{s^\tau | s^t} \pi(s^\tau | s^t) \frac{p(s^\tau)}{R^{\tau - t}}$$

Intuitively, the ML manages risk by enforcing a strict credit limit: the agent's debt can never exceed a specific fraction k of their total future wealth.

Assumption 3.4: Borrowing Limit

The ML will define a parameter $k \in [0, 1]$ such that:

$$D(s^t) \leq kW(s^t), \quad \forall s^t.$$

In default, the agent can no longer borrow from the ML, but can still save by purchasing cash-in-advance contracts (Arrow securities).

Definition 3.5: Cash-In-Advance

Cash-in-advance contracts allow the agent to save by purchasing state-contingent claims on future payoffs. The agent can choose a saving strategy defined by:

$$a(s^t) = \sum_{s^{t+1} | s^t} \pi(s^{t+1} | s^t) \frac{g(s^{t+1})}{R}$$

where $g(s^{t+1}) \geq 0$ is the state-contingent return on the asset tomorrow.

The intuition if the cash-in-advance contracts is that the agent acts as their own actuary. $a(s^t)$ is the total premium (cash) the agent pays today. In exchange, the market guarantees a non-negative payout $g(s^{t+1})$ tomorrow, exactly tailored to the agent's desired self-insurance needs.

Theorem 3.6: Bulow-Rogoff

In any equilibrium, the debt must be non-positive ($D(s^t) \leq 0$) for all s^t, the borrowing limit $k = 0$, and the agent cannot borrow at all.

Here, we motivate the proof by intuition before going into the formal proof.

Suppose there exists a node s^t where the debt satisfies two conditions:

1. $D(s^t) \leq kW(s^t)$ (The global borrowing limit is respected)
2. $D(s^t) > k\left(W(s^t) - y(s^t)\right)$ (The debt exceeds future collateral)

The intuition for the second condition is that, the current debt is so heavy that it exceeds the allowable borrowing limit on *strictly future* wealth ($W - y$). Consequently, the agent is forced to make a painful net transfer out of today's pocket ($y(s^t)$) to service the debt. This is the optimal moment to default.

We consider the case when the agent defaults at this specific node s^t and starts the following savings sequence for all $\tau \geq t$:

$$a(s^\tau) = p(s^\tau) + k\left(W(s^\tau) - y(s^\tau)\right) - D(s^\tau),$$

where k is assumed to be $k > 0$, since we try to prove by contradiction that k must be zero in equilibrium.

Here is the intuition for this saving sequence: the agent redirects the payment $p(s^\tau)$ they *would* have made to the ML into a private savings account. They adjust this amount by taking the maximum theoretical future borrowing limit ($k(W - y)$) and subtracting the actual future debt burden (D). In effect, they use the financial market to mimic the ML's balance sheet.

But we are still left with two critical questions:

- **Feasibility:** Is this asset (defined by the savings sequence) legally available in the market? In other words, can the agent actually purchase this asset to implement the savings strategy?
- **Optimality:** Does this asset make the agent strictly better off than repaying the debt? In other words, is defaulting and following this savings strategy strictly more profitable than repaying the debt?

To prove this asset is legally available in the market, we must show that its future payoff $g(s^{\tau+1})$ is non-negative.

First, we write down the recursive definitions of W and D:

$$W(s^\tau) = y(s^\tau) + \sum_{s^{\tau+1}|s^\tau} \pi(s^{\tau+1}|s^\tau)\frac{W(s^{\tau+1})}{R} \implies W(s^\tau) - y(s^\tau) = \sum_{s^{\tau+1}|s^\tau} \pi(s^{\tau+1}|s^\tau)\frac{W(s^{\tau+1})}{R}$$

$$D(s^\tau) = p(s^\tau) + \sum_{s^{\tau+1}|s^\tau} \pi(s^{\tau+1}|s^\tau)\frac{D(s^{\tau+1})}{R} \implies D(s^\tau) - p(s^\tau) = \sum_{s^{\tau+1}|s^\tau} \pi(s^{\tau+1}|s^\tau)\frac{D(s^{\tau+1})}{R}$$

Substitute these recursive forms back into our savings sequence $a(s^\tau)$:

$$\begin{aligned}
a(s^\tau) &= k\left(W(s^\tau) - y(s^\tau)\right) - \left(D(s^\tau) - p(s^\tau)\right) \\
&= k \sum_{s^{\tau+1}|s^\tau} \pi(s^{\tau+1}|s^\tau)\frac{W(s^{\tau+1})}{R} - \sum_{s^{\tau+1}|s^\tau} \pi(s^{\tau+1}|s^\tau)\frac{D(s^{\tau+1})}{R} \\
&= \sum_{s^{\tau+1}|s^\tau} \pi(s^{\tau+1}|s^\tau)\frac{kW(s^{\tau+1}) - D(s^{\tau+1})}{R}
\end{aligned}$$

By matching this with the standard cash-in-advance pricing formula ($a(s^\tau) = \sum \pi \frac{g}{R}$), we identify the state-contingent payoff tomorrow as:

$$g(s^{\tau+1}) = kW(s^{\tau+1}) - D(s^{\tau+1}).$$

Since the ML's contract globally enforces $D(s^{\tau+1}) \leq kW(s^{\tau+1})$, it strictly follows that $g(s^{\tau+1}) \geq 0$ for all possible future states. Therefore, the proposed asset is a feasible cash-in-advance contract.

Then, we compare the agent's consumption payoffs under repayment vs. default.

- At the moment of default (Today, s^t):
  - Payoff in repayment: $y(s^t) - p(s^t)$
  - Payoff in default:

$$\begin{aligned}
y(s^t) - a(s^t) &= y(s^t) - \left[p(s^t) + k\left(W(s^t) - y(s^t)\right) - D(s^t)\right] \\
&= y(s^t) - p(s^t) + \underbrace{\left[D(s^t) - k\left(W(s^t) - y(s^t)\right)\right]}_{>0 \text{ (due to our choice of the tipping point)}}
\end{aligned}$$

  As a result, the agent consumes strictly more today by defaulting.

- At any future node (Tomorrow and beyond, s^τ for $\tau > t$):
  - Payoff in repayment: $y(s^\tau) - p(s^\tau)$
  - Payoff in default (= endowment - cost of new asset + payoff from yesterday's asset):

$$\begin{aligned}
&y(s^\tau) - a(s^\tau) + g(s^\tau) \\
&= y(s^\tau) - [p(s^\tau) + k(W(s^\tau) - y(s^\tau)) - D(s^\tau)] + [kW(s^\tau) - D(s^\tau)] \\
&= y(s^\tau) - p(s^\tau) - kW(s^\tau) + ky(s^\tau) + D(s^\tau) + kW(s^\tau) - D(s^\tau) \\
&= (1+k)y(s^\tau) - p(s^\tau)
\end{aligned}$$

  As a result, since the endowment $y > 0$, if $k > 0$, then $(1+k)y - p > y - p$. The agent consumes strictly more (ky extra) in every single future period.

In conclusion, if $k > 0$, the agent will inevitably reach node s^t, choose to default, and be strictly better off in every period thereafter. Anticipating this, the ML will never lend. Consequently, the only equilibrium is $k \leq 0 \implies k = 0$, meaning debt $D(s^t) \leq 0$, and the agent cannot borrow.

Remark (Existence of Tipping Point and Discussion of k).

The proof above relies on the existence of a **tipping point** s^* that satisfies the two conditions. Existence is guaranteed by the following lemma.

Lemma 3.7: Existence of the Tipping Point

Suppose the Money Lender (ML) offers a contract where the supremum of the debt-to-wealth ratio is strictly positive. Let $k = \sup_{s^t} \frac{D(s^t)}{W(s^t)} > 0$. Then, there exists at least one node s^* in the contract history that satisfies:

1. $D(s^*) \leq kW(s^*)$
2. $D(s^*) > k\left(W(s^*) - y(s^*)\right)$

Proof for Lemma

By the definition of the supremum k, the first condition $D(s^t) \leq kW(s^t)$ holds globally for all s^t.

To prove the second condition, we first establish a strictly positive lower bound for the ratio of the flow endowment to the total stock wealth, $\frac{y(s^t)}{W(s^t)}$. Assuming a finite Markov state space where endowments are strictly positive, there exists a minimum possible endowment $y_{\min} > 0$ and a maximum possible total wealth $W_{\max} < \infty$ (since the discount rate $R > 1$). Therefore, for any node s^t:

$$\frac{y(s^t)}{W(s^t)} \geq \frac{y_{\min}}{W_{\max}} \equiv \delta > 0$$

Now, choose an arbitrarily small positive number ϵ such that $\epsilon = \frac{1}{2}k\delta > 0$. By the definition of the supremum k, there must exist at least one node s^* such that the actual debt-to-wealth ratio is strictly within ϵ of the supremum:

$$\frac{D(s^*)}{W(s^*)} > k - \epsilon \implies D(s^*) > kW(s^*) - \epsilon W(s^*)$$

Substitute our specific choice of ϵ into this inequality:

$$\begin{aligned} D(s^*) &> kW(s^*) - \left(\frac{1}{2}k\delta\right) W(s^*) \\ &> kW(s^*) - k\left(\frac{y(s^*)}{W(s^*)}\right) W(s^*) \\ &= k\left(W(s^*) - y(s^*)\right). \end{aligned}$$

Thus, the node s^* satisfying both conditions is mathematically guaranteed to exist. ∎

For narrative simplicity, we often state that the Money Lender (ML) specifies a global borrowing limit k such that $D(s^t) \leq kW(s^t)$. However, mathematically, it is crucial that the k used in this proof is defined as the *supremum* of the actual debt-to-wealth ratio, rather than an arbitrary constant.

If we were to use a loosely defined nominal rule k_{rule}, and the actual borrowing behavior is significantly more conservative (i.e., $k_{\text{rule}} \gg \sup \frac{D}{W}$), the strict inequality

> $D(s^t) > k_{\text{rule}}\,(W(s^t) - y(s^t))$ might never be satisfied. In such a case, the debt would never mathematically "squeeze" the current endowment enough to trigger the tipping point. Therefore, constructing the shadow asset around the precise, tight supremum k is the exact mechanism that guarantees the existence of the default node.

3.3 Two-Sided Lack of Commitment

Assume:

- Two agents:
  - identical preferences with risk aversion;
  - lack of commitment.
- Endowments:
  $$y_A = y_s, \quad y_B = 1 - y_s, \quad y_s \in [0, 1],$$
  where y_s is i.i.d. with each state s occurring with probability $\pi(s)$.
  - The feasibility constraint is then given by
  $$c_A(s^t) + c_B(s^t) = 1, \quad \forall t, s^t.$$
- Lending and borrowing can only be done through the reallocation of consumption between agents. If any agent defaults on the contract, then both agents will be in autarky and consume their own endowment forever.
  $$V_A^{\text{aut}} = \sum_t \beta^t \sum_s \pi(s) u(y_s) = \frac{\sum_s \pi_s u(y_s)}{1 - \beta},$$
  $$V_B^{\text{aut}} = \sum_t \beta^t \sum_s \pi(s) u(1 - y_s) = \frac{\sum_s \pi_s u(1 - y_s)}{1 - \beta}.$$

An allocation $\{c_A(s^t), c_B(s^t)\}_{t=0}^{\infty}$ is sustainable if neither agent has incentives to default on the contract at any point in time. That is, for all t and s^t, we have

$$u(c^A(s^t)) + \mathbb{E}_t \left[\sum_j \beta^j u(c^A(s^{t+j})) \right] \geq u(y(s^t)) + \beta V_A^{\text{aut}},$$
$$u(c^B(s^t)) + \mathbb{E}_t \left[\sum_j \beta^j u(c^B(s^{t+j})) \right] \geq u(1 - y(s^t)) + \beta V_B^{\text{aut}}.$$

We characterize the Pareto frontier of feasible and sustainable allocation by the function $Q(v)$, which is defined as the maximum utility that agent B can get given that agent A gets at least v more utility than their autarky utility (we call it *difference in utilities* later):

$$\begin{aligned} Q(v) = \max_{c(s^t)} & \sum_t \sum_{s^t} \beta \left[u(1-c(s^t)) - u(1-y(s^t))\right] \\ \text{s.t.} \quad & \sum_t \sum_{s^t} \beta \left[u(c(s^t)) - u(y(s^t))\right] \geq v, \\ & c(s^t) \in \Gamma. \end{aligned}$$

where Γ is the set of feasible and sustainable allocations.

Obviously, it has to be true that $v \geq 0$ and $Q(v) \geq 0$ for all $v \geq 0$.

If $\underline{v} = 0$, then the present value of A's consumption utility is promised to be just at least as high as A's autarky utility, so $v_{\min} = \underline{v} = 0$. If $Q(\bar{v}) = 0$, then the present value of B's consumption utility is promised to be just at least as high as B's autarky utility, so $v_{\max} = \bar{v}$.

The recursive problem can be written as

$$\begin{aligned} Q(\Delta, s) = \max_{c, \Delta(s')} & \; u(1-c) - u(1-y_s) + \beta \sum_{s'} \pi(s') Q(\Delta(s'), s') \\ \text{s.t.} \quad & u(c) - u(y_s) + \beta \sum_{s'} \pi(s') \Delta(s') \geq \Delta, \\ & \Delta(s') \geq 0, \quad \forall s', \\ & Q(\Delta(s'), s') \geq 0, \quad \forall s', \\ & c \in [0, 1]. \end{aligned}$$

Here Δ is the promise in terms of the difference in utilities for agent A (i.e., the present value of A's consumption utility minus A's autarky utility). Note that the first constraint is the promise-keeping constraint, and the second and third constraints are the participation constraints for agent A and agent B, respectively.

Remark.

- We can show that $\Delta(s) \in \left[0, \bar{\Delta}(s)\right]$ for all s, where $\bar{\Delta}(s)$ is such that $Q(\bar{\Delta}(s), s) = 0$.
- Q inherits the properties of u: Q is decreasing in Δ, strictly concave in Δ, and continuously differentiable in Δ.

Let μ be the Lagrange multiplier for the promise-keeping constraint, $\lambda_{s'} \beta \pi_{s'}$ be the Lagrange multiplier for the participation constraint of agent A, and $\theta_{s'} \beta \pi_{s'}$ be the Lagrange multiplier for the participation constraint of agent B. The FOC is given by

- w.r.t. c:
$$-u'(1-c) + \mu u'(c) = 0.$$
- w.r.t. $\Delta(s')$:
$$\beta \pi_{s'} Q'(\Delta(s'), s') + \mu \beta \pi_{s'} + \lambda_{s'} \beta \pi_{s'} + \theta_{s'} \beta \pi_{s'} Q'(\Delta(s'), s') = 0.$$

The envelope condition is given by

$$Q'(\Delta, s) = -\mu.$$

From the FOC and the envelope condition, we can obtain the following results:

$$\begin{aligned} Q'(\Delta, s) &= -\frac{u'(1-c)}{u'(c)}, \\ Q'(\Delta, s) &= (1+\theta_{s'})Q'(\Delta(s'), s') + \lambda_{s'}. \end{aligned}$$

This implies that the consumption c is increasing in Δ. We have already argued that $\Delta(s) \in \big[0, \bar{\Delta}(s)\big]$ for all s. Therefore, the consumption c will fall in the range $[\underline{c}_s, \bar{c}_s]$ for all s, where

$$\begin{aligned} Q'(0, s) &= -\frac{u'(1-\underline{c}_s)}{u'(\bar{c}_s)}, \quad (c = \underline{c}_s, \Delta = 0), \\ Q'(\bar{\Delta}(s), s) &= -\frac{u'(1-\bar{c}_s)}{u'(\underline{c}_s)}, \quad (c = \bar{c}_s, \Delta = \bar{\Delta}(s)). \end{aligned}$$

Furthermore, we can rewrite the FOC w.r.t. c as

$$c = g(Q'(\Delta, s)) = g\left(-\frac{u'(1-c)}{u'(c)}\right),$$

where g^{-1} is decreasing and negative.

With this, the FOC w.r.t. $\Delta(s')$ can be rewritten as

$$g^{-1}(c_s) = (1+\theta_{s'})g^{-1}(c_{s'}) + \lambda_{s'}.$$

Here we analyze the problem by considering four cases:

- No participation constraint binds. Then $\lambda_{s'} = 0$ and $\theta_{s'} = 0$ for all s'. This implies that $Q'(\Delta, s) = Q'(\Delta(s'), s')$ and $g^{-1}(c_s) = g^{-1}(c_{s'})$ for all s and s', so the consumption is constant across states:

$$c_s = c_{s'}, \quad \Delta = \Delta(s').$$

- Only participation constraint of agent A binds. Then $\lambda_{s'} > 0$ and $\theta_{s'} = 0$ for some s'. Also by construction, $\Delta(s') = 0$ and $c_{s'} = \underline{c}_{s'}$ in this case. So we have

$$\begin{aligned} & g^{-1}(c_s) = g^{-1}(c_{s'}) + \lambda_{s'} \\ \Longrightarrow\ & g^{-1}(c_s) > g^{-1}(c_{s'}) \\ \Longrightarrow\ & c_s < c_{s'}. \end{aligned}$$

- Only participation constraint of agent B binds. Then $\lambda_{s'} = 0$ and $\theta_{s'} > 0$ for some s'. Also by construction, $\Delta(s') = \bar{\Delta}(s')$, $Q(\Delta(s'), s') = 0$, and $c_{s'} = \bar{c}_{s'}$ in this case. So we have

$$\begin{aligned} & g^{-1}(c_s) = (1+\theta_{s'})g^{-1}(c_{s'}) \\ \Longrightarrow\ & g^{-1}(c_s) < g^{-1}(c_{s'}) \\ \Longrightarrow\ & c_s > c_{s'} \\ \Longrightarrow\ & 1 - c_s < 1 - c_{s'}. \end{aligned}$$

- Both participation constraints bind.

  This case is not feasible as long as some risk-sharing is possible. If both participation constraints bind simultaneously in some state s', it implies that $\Delta(s') = 0$ (Agent A gets exactly their autarky value) and $Q(\Delta(s'), s') = 0$ (Agent B gets exactly their autarky value). Mathematically, this means $Q(0, s') = 0$.

  However, this contradicts the assumption that some risk-sharing is possible. If risk-sharing is sustainable, the contract generates a strictly positive total surplus relative to autarky due to the agents' risk aversion.

  Since $Q(\Delta, s')$ defines the *Pareto frontier* (the maximum possible surplus for B given A's surplus), if A is pushed down to their autarky outside option ($\Delta = 0$), Agent B must capture *all* the strictly positive risk-sharing surplus. Therefore, it must be true that $Q(0, s') > 0$.

  In geometric terms, the autarky point $(0, 0)$ lies strictly *inside* the feasible payoff set. The Pareto frontier Q strictly bounds this set away from the origin. The only scenario where $Q(0, s') = 0$ holds is when the commitment friction is so severe that the risk-sharing market completely collapses, making autarky the only sustainable allocation.

In conclusion,

$$c_{s'} = \begin{cases} c_s & \text{if } c_s \in [\underline{c}_{s'}, \bar{c}_{s'}] \quad \text{(no participation constraint binds)} \\ \underline{c}_{s'} & \text{if } c_s < \underline{c}_{s'}, \Delta(s') = 0 \quad \text{(only participation constraint of agent } A \text{ binds)} \\ \bar{c}_{s'} & \text{if } c_s > \bar{c}_{s'}, Q(\Delta(s'), s') = 0 \quad \text{(only participation constraint of agent } B \text{ binds)} \end{cases}$$

Remark.

- **Intuition**: When the participation constraint of agent A binds, the planner raises the consumption of agent A so that they are indifferent between autarky and the contract, given their high endowment today. The same logic applies to agent B when the participation constraint of agent B binds.

- Note that although agent A is getting the minimum consumption level $\underline{c}_{s'}$ when the participation constraint of agent A binds, it does not necessarily mean that agent A is worse off than in autarky. Actually agent A is made (weakly) better off than in the previous state because the planner raises their consumption level to keep them in the contract. (The "worst" in a good state is still better than some consumption level in a bad state.)

Claim

We can obtain a stationary distribution of consumption if

- consumption intervals do not contain each other. That is,

$$y_1 > y_2 \implies \bar{c}_1 > \bar{c}_2, \quad \underline{c}_1 > \underline{c}_2.$$

- endowments are contained in the consumption intervals. That is,

$$y_s \in [\underline{c}_s, \bar{c}_s], \quad \forall s.$$

- consumption intervals are non-degenerate. That is,

$$\underline{c}_s < \bar{c}_s, \quad \forall s.$$

Proof for Claim.

Claim

$$Q(\Delta + u(y_2) - u(y_1), s_1) = Q(\Delta, s_2) + u(1 - y_2) - u(1 - y_1).$$

Proof for Claim.

We prove this by definition.

$$\begin{aligned} Q(\Delta + u(y_2) - u(y_1), s_1) = \max u(1-c) - u(1-y_1) + \beta \sum_{s'} \pi(s') Q(\Delta(s'), s') \\ \text{s.t. } u(c) - u(y_1) + \beta \sum_{s'} \pi(s')\Delta(s') \geq \Delta + u(y_2) - u(y_1). \end{aligned}$$

$$\begin{aligned} Q(\Delta, s_2) + u(1-y_2) - u(1-y_1) = \max u(1-c) - u(1-y_2) + \beta \sum_{s'} \pi(s') Q(\Delta(s'), s') \\ \text{s.t. } u(c) - u(y_2) + \beta \sum_{s'} \pi(s')\Delta(s') \geq \Delta. \end{aligned}$$

Hence, the equality holds. ∎

Let $y_1 > y_2$. We evaluate at $\Delta = \bar{\Delta}(s_2) := \bar{\Delta}_2$. By the just-proven equality, we have

$$Q(\bar{\Delta}_2 + u(y_2) - u(y_1), s_1) = Q(\bar{\Delta}_2, s_2) + u(1 - y_2) - u(1 - y_1).$$

Note that $Q(\bar{\Delta}_2, s_2) = 0$ by definition, and $u(1-y_2) - u(1-y_1) > 0$ since $y_1 > y_2$ and u is increasing. Therefore, we have $Q(\bar{\Delta}_2) > 0$. Also note that $0 = Q(\bar{\Delta}_1, s_1)$, so we have

$$Q(\bar{\Delta}_2 + u(y_2) - u(y_1), s_1) > Q(\bar{\Delta}_1, s_1).$$

Since Q is decreasing in Δ, this implies

$$\bar{\Delta}_2 + u(y_2) - u(y_1) < \bar{\Delta}_1.$$

Since Q is strictly concave, Q' is strictly decreasing in Δ. Therefore, we have

$$\begin{aligned} Q'(\bar{\Delta}_1, s_1) &< Q'(\bar{\Delta}_2 + u(y_2) - u(y_1), s_1) \\ &= Q'(\bar{\Delta}_2, s_2), \end{aligned}$$

where the equality holds by the just-proven equality (and taking derivatives).

Note that we previously defined $c = g(Q'(\Delta, s))$, and we have argued that g is strictly increasing. Finally, we have

$$\bar{c}_1 > \bar{c}_2.$$

Remark (Endogenously Incomplete Markets: Where We Have Arrived).

The three settings developed in this chapter—one-sided lack of commitment, Bulow–Rogoff with re-access to savings markets, and two-sided lack of commitment—share a common methodological signature: rather than *assuming* that the asset market is incomplete (as in the previous chapter), they let the asset market *become* incomplete as the equilibrium response to a participation friction. Three lessons survive across the three models.

- **The participation constraint is a state variable in disguise.** Whether expressed as a sequence of $u(c) + \beta w \geq u(y) + \beta V^{\text{aut}}$ inequalities (one-sided), as a wealth-to-debt ratio bound (Bulow–Rogoff), or as a pair of value-function constraints (two-sided), the shadow value of relaxing the participation constraint enters the dynamics directly. Promised utility v (or Δ) is the natural recursive state because it summarizes everything the planner needs to remember.

- **Risk sharing is partial and history-dependent.** Optimal contracts implement *conditional* consumption smoothing: consumption is constant within "intervals" of states (where no participation constraint binds) and adjusts only when a constraint binds. The result is a non-trivial stationary distribution of consumption—qualitatively similar to what exogenously incomplete markets produce, but micro-founded by the friction.

- **What outside option is available matters enormously.** Bulow–Rogoff is the cleanest illustration: relative to one-sided LoC, the only change is that the defaulting agent is allowed to save in complete markets afterward, and the entire borrowing market collapses. The lesson generalizes: any structural model of credit-market frictions stands or falls with its specification of what the agent can do after default.

These insights motivate the heterogeneous-agent macro models we turn to in later chapters: the cross-sectional wealth distribution is not just a residual of model-driven heterogeneity in shocks, but the equilibrium outcome of microfounded participation and information frictions.

Remark (Chapter Summary).

- **Endogenously incomplete markets.** Rather than restricting the asset menu by fiat, this chapter derives incompleteness as the equilibrium response to a participation friction. The set of sustainable contracts is shaped by what each party can do upon walking away.

- **Promised utility as the recursive state.** In the one-sided LoC model, v summarizes the lender's outstanding obligation. In the two-sided model, Δ (utility surplus over autarky) plays the same role. The Pareto frontiers $P(v)$ and $Q(\Delta, s)$ encode the constrained-efficient allocations.

- **Contracts smooth conditionally.** Consumption is constant within "intervals" of states (no participation constraint binds) and adjusts only when one binds. The result is an endogenous, history-dependent stationary consumption distribution—qualitatively similar to the exogenously incomplete case but micro-founded.

- **Bulow–Rogoff: the outside option is everything.** Allow defaulters to save in complete cash-in-advance markets, and the entire borrowing market collapses ($k = 0$ in equilibrium). The lesson: any structural model of credit-market frictions stands or falls with its specification of post-default options.

- **Two-sided LoC stationary distribution.** Existence of a non-degenerate stationary distribution of consumption requires the consumption intervals not to nest—i.e., risk sharing is partial but persistent.

- **Connection to Part II.** The wealth distributions that drive HANK and Aiyagari-style models can be interpreted as reduced forms of these participation-friction equilibria, justifying their economic content.

Part II

Growth, Business Cycles, and Quantitative Macroeconomics

Lectures by Kai-Jie Wu

Chapter 4

Growth and Development Accounting

Remark (Notation in This Chapter).

Symbol	Meaning
$Y_{it}, K_{it}, H_{it}, L_{it}$	Output, physical capital, total human capital, raw labor force in country i at t
y, k, h	Per-worker counterparts
α	Capital share / output elasticity in Cobb–Douglas $y = Ak^\alpha h^{1-\alpha}$
$j \in \{1, \ldots, J\}$	Schooling group index
s_j	Years of schooling for group j
l_j	Population share of schooling group j
h_j	Productivity weight for schooling group j (relative to $h_1 \equiv 1$)
w_j	Wage rate for schooling group j
m	Mincer return to one additional year of schooling
$\hat{h}_{it}$	Empirical estimator for human capital using Mincer wage premium
g_x	Average growth rate of variable x over a sample period
s_k, s_h, s_A	Decomposition shares of k, h, TFP in income variation / growth
X_i	Composite of observable inputs: $X_i \equiv \alpha \ln k_i + (1-\alpha) \ln h_i$

This chapter transitions from theoretical dynamics to empirical measurement through an exercise known as **Accounting** (i.e., variance and growth decomposition). The fundamental goal of this chapter is to open the "black box" of aggregate output and attribute it to observable factors versus unobservable productivity. We ask two distinct quantitative questions:

1. **Development Accounting (Cross-Sectional):** Why are some countries vastly richer than others at a given point in time? How much of this income gap is due to differences in physical and human capital, versus differences in Total Factor Productivity (TFP)?

2. **Growth Accounting (Time-Series):** Why does a specific country grow over time?

How much of its historical growth rate can be attributed to the accumulation of capital, versus the growth of TFP?

Typically, empirical evidence suggests that physical capital accumulation accounts for 30–40% of growth, human capital for about 10% (though debatable), and the remaining roughly 50% is attributed to improving productivity.

4.1 The Production Framework: From General to Specific

Assume we observe a panel dataset of economies $i \in I$ over time periods $t \in T$:

$$\{Y_{it}, K_{it}, H_{it}, L_{it}\}_{i \in I, t \in T}$$

where Y is total output, K is physical capital, H is total human capital (efficiency units of labor, defined later), and L is the raw labor force.

The aggregate production function is given by:

$$Y_{it} = A_{it} F(K_{it}, H_{it})$$

where $F(\cdot)$ is assumed to exhibit *constant returns to scale (CRS)*.

Here, it is crucial to distinguish between the physical properties of the inputs:

- **Rivalrous Inputs (K_{it}, H_{it}):** A good is defined as *strictly rivalrous* if its use or consumption by one economic agent fundamentally precludes its simultaneous use by another. Physical capital and human capital are structurally rivalrous because they are strictly bound by physical space and time constraints. For instance, a single machine (K) cannot be simultaneously operated by multiple workers without severe congestion. Similarly, human capital (H) is embodied within the worker; an hour of a worker's skilled labor deployed in one factory physically cannot be concurrently deployed in another.

- **Non-Rivalrous Productivity (A_{it}):** The Total Factor Productivity (TFP) term A_{it} captures non-rivalrous goods—such as knowledge, blueprints, software, and institutional quality. An idea can be shared simultaneously across all workers and machines without being depleted.

Because $F(\cdot)$ is CRS, we can divide both sides of the equation by the raw labor force L_{it} to express the economy in per-worker terms. Crucially, because A_{it} is non-rivalrous, this transformation does not dilute it; every worker can fully access exactly the same level of technology[1].

$$y_{it} = A_{it} F\left(\frac{K_{it}}{L_{it}}, \frac{H_{it}}{L_{it}}\right) \equiv A_{it} f(k_{it}, h_{it})$$

where y_{it} is output per worker, k_{it} is physical capital per worker, and h_{it} is human capital per worker.

While the general function $f(\cdot)$ establishes the abstract theoretical mechanics, performing an actual quantitative accounting exercise requires concrete numbers. Specifically, we need to know *by exactly what percentage* output changes when an input changes—this is the output elasticity. To measure and decompose variances in the real-world data, we must parameterize these output elasticities into specific, measurable constants. The universal standard in the macro literature is to adopt the Cobb-Douglas functional form, which conveniently assumes a constant capital share (elasticity) $\alpha \in (0, 1)$:

$$y_{it} = A_{it} k_{it}^{\alpha} h_{it}^{1-\alpha}$$

Taking the natural logarithm of both sides, we obtain the fundamental log-linearized accounting identity:

$$\ln y_{it} = \ln A_{it} + \alpha \ln k_{it} + (1 - \alpha) \ln h_{it}$$

This linear additive form is the starting point for all variance and growth decompositions. The fundamental question of macro accounting is: How much of the variation in $\ln y_{it}$ is driven by the observable inputs (k_{it}, h_{it}) versus the unobservable TFP residual (A_{it})? Before answering it, the next section addresses two prerequisite measurement problems: how to pin down α, and how to measure human capital h_{it}.

4.2 Measurement: α and Human Capital

4.2.1 How to Pin Down α?

To conduct the accounting exercise, we need the value of α. A naive approach is to simply run an OLS regression of $\ln y_i$ on $\ln k_i$ and $\ln h_i$ to extract the coefficient. However, this suffers from severe *endogeneity (omitted variable bias)*. The error term $\ln A_i$ is likely to be positively correlated with the regressors (e.g., higher tech countries attract more capital k_i), leading OLS to overestimate α.

There are typically two ways to cope with this endogeneity problem:

1. **Econometrics:** Find an IV that is orthogonal to $\ln A_i$ and run a 2SLS regression. However, this is often difficult in practice, and the results can be sensitive to the choice of instrument.

[1]Algebraically, because $F(\cdot)$ is CRS, dividing both sides by L_{it} allows us to pass $1/L_{it}$ inside the function F, yielding $y_{it} = A_{it} f(k_{it}, h_{it})$. However, this mathematical legality must be strictly backed by the underlying physical properties in economics.

Imagine if A_{it} were not abstract technology, but rather a pile of "coal" (a rivalrous good). The original production function would have to be written as $Y_{it} = F(K_{it}, H_{it}, A_{it})$. When converting this economy to per-worker terms, the coal allocated to each worker would inevitably be diluted to $a_{it} = A_{it}/L_{it}$.

The only reason we are allowed to place A_{it} *outside* the function F as a global multiplier—and completely exempt it from being divided by L_{it} during the per-worker transformation—is because we have predefined technology as strictly non-rivalrous. Regardless of how much the population L_{it} expands, the "calculus formula" (A_{it}) available to each worker remains the complete A_{it}, rather than a diluted A_{it}/L_{it}.

2. **Economic Theory (the standard approach):** Model the bias away using microeconomics.

Assuming perfectly competitive markets[2], the representative firm's profit maximization problem is:

$$\max_{k_i, h_i} A_i k_i^{\alpha} h_i^{1-\alpha} - r k_i - w h_i$$

Taking the FOC with respect to k_i gives the marginal product of capital:

$$\alpha A_i k_i^{\alpha-1} h_i^{1-\alpha} = r \implies \alpha \frac{y_i}{k_i} = r$$

Rearranging this yields:

$$\alpha = \frac{r k_i}{y_i}$$

This shows that α is exactly equal to the *capital income share* in the national accounts. We can simply look up this share in the data without running any biased regressions.

4.2.2 How to Measure Human Capital h_{it}?

Following the seminal work of Hall and Jones (1999), we construct aggregate human capital by linking macroeconomic aggregates to microeconomic wage data.

Suppose a country has heterogeneous workers divided into schooling groups $j \in \{1, 2, \cdots, J\}$ (e.g., No schooling, Elementary, High school, College). Each group has s_j years of schooling, and makes up a share l_j of the population. Let group 1 be the baseline uneducated workers ($h_1 = 1$).

Aggregate human capital H is the sum of efficiency units of labor:

$$H = \sum_{j=1}^{J} h_j L_j$$

This equation converts a heterogeneous workforce into a single, homogeneous macroeconomic measure called "efficiency units of labor." Specifically, L_j is the absolute number of workers in schooling group j, and h_j captures their specific human capital level (or relative productivity weight). By normalizing $h_1 = 1$ for uneducated workers, H essentially measures the total labor input expressed in terms of "uneducated-worker equivalents."

The firm's problem is to choose capital and the amount of labor from each group to maximize profit:

$$\max_{K, \{L_j\}} A K^{\alpha} \left(\sum_{j=1}^{J} h_j L_j \right)^{1-\alpha} - rK - \sum_{j=1}^{J} w_j L_j$$

The FOC with respect to labor type j equates the marginal product of labor to its wage:

$$w_j = (1-\alpha) A K^{\alpha} H^{-\alpha} \cdot h_j = \left((1-\alpha) \frac{Y}{H} \right) h_j$$

[2] In standard microeconomic theory, the assumption of a perfectly competitive market implies that the representative firm is a strict *price taker*. Consequently, the factor prices—the rental rate of capital (r) and the wage rate (w)—are exogenous and fixed from the firm's perspective. This allows us to equate the marginal product of capital directly to r without modeling the firm's impact on market prices.

By dividing the FOC of group j by the FOC of the uneducated group 1, we obtain the relative wage ratio:

$$\frac{w_j}{w_1} = \frac{h_j}{h_1} = h_j$$

This shows that *relative wages perfectly reflect relative human capital (productivity)*. Thus, the aggregate and per-capita human capital measures H and h can be expressed as[3]:

$$\hat{H} = \sum_{j=1}^{J} \left(\frac{w_j}{w_1} \right) L_j$$

$$\hat{h} = \sum_{j=1}^{J} \left(\frac{w_j}{w_1} \right) l_j$$

4.2.3 The Mincer Equation

How do we find the wage ratio w_j/w_1? We rely on the Mincer (1974) wage regression from labor economics:

$$\ln w_i = ms_i + \beta X_i + \varepsilon_i$$

Here, m is the *Mincer return*, capturing the percentage wage gain from one additional year of schooling (s_i). X_i is a vector of controls. This equation implies:

$$\ln w_j - \ln w_1 = m(s_j - s_1) \implies \frac{w_j}{w_1} = e^{m(s_j - s_1)}$$

Substituting the Mincerian wage premium back into our human capital index, we get the final measurement formula:

$$\hat{h}_{it} = \sum_{j=1}^{J} e^{m_{it}(s_{jit} - s_{1it})} l_{jit}$$

where

- i: denotes the specific country (the cross-sectional dimension).
- t: denotes the specific year (the time-series dimension).
- j: denotes the specific schooling group within that country and year.

To compute this across countries and time, we combine two datasets:

- $\{l_{jit}, s_{jit}\}$: Population shares and years of schooling from **Barro and Lee (2013)**.
- $\{m_{it}\}$: Estimates of Minceerian returns to schooling from **Psacharopoulos and Patrinos (2018)**.

[3] The "hat" (ˆ) explicitly denotes that this is an *empirical estimate* or proxy. The true human capital level h is intrinsically unobservable. However, by leveraging the microeconomic equilibrium condition that wage ratios equal productivity ratios, we can construct the estimator $\hat{h}$ using purely observable market data (relative wages and population shares).

4.3 Development Accounting

Fix a specific year t, yielding a cross-sectional dataset $\{y_i, k_i, h_i\}_{i\in I}$. To simplify notation, we aggregate the observable inputs into a single "factor composite" X_i:

$$X_i \equiv \alpha \ln k_i + (1-\alpha)\ln h_i$$

Thus, the income variation is perfectly decomposed into observable factors and unobservable TFP:

$$\ln y_i = X_i + \ln A_i$$

- **Method 1: The Simple Variance Ratio**
  A naive approach measures the contributions as:

  $$s_k^{(1)} = \frac{\text{Var}[\alpha \ln k_i]}{\text{Var}[\ln y_i]}, \quad s_h^{(1)} = \frac{\text{Var}[(1-\alpha)\ln h_i]}{\text{Var}[\ln y_i]}$$

  Limitations: This method has two major statistical flaws. First, variances cannot capture the direction of the relationship. Second, because $\text{Var}[\ln y_i] = \text{Var}[X_i] + \text{Var}[\ln A_i] + 2\,\text{Cov}[X_i, \ln A_i]$, the shares will not sum to 100% unless the covariance between observable factors and TFP is exactly zero (which is strongly rejected in the data).

- **Method 2: Covariance-Based Decomposition**
  To address these limitations, the modern literature (e.g., Klenow and Rodríguez-Clare, 1997) standardly adopts a covariance-based approach:

  $$s_k = \frac{\text{Cov}[\alpha \ln k_i, \ln y_i]}{\text{Var}[\ln y_i]}, \quad s_h = \frac{\text{Cov}[(1-\alpha)\ln h_i, \ln y_i]}{\text{Var}[\ln y_i]}$$

  Intuition: Notice that s_k is exactly the slope coefficient β_k from a bivariate OLS regression of $\alpha \ln k_i$ on $\ln y_i$. Furthermore, because covariance is a linear operator, this method splits the joint covariance term linearly and guarantees exact additivity: $s_k + s_h + s_A = 1$ (if we in addition define $s_A = \frac{\text{Cov}[\ln A_i, \ln y_i]}{\text{Var}[\ln y_i]}$).[4]

4.4 Growth Accounting

While development accounting focuses on the cross-country differences in income *levels* at a given point in time, *growth accounting* aims to decompose the changes in income *over time*.

Suppose we have successfully pinned down the capital share α and observed the panel dataset:

$$\{y_{it}, k_{it}, h_{it}\}_{i\in I, t\in T}$$

[4] By the definition of variance, we know that $\text{Var}[\ln y_i] = \text{Cov}[\ln y_i, \ln y_i]$. Substituting the income identity $\ln y_i = X_i + \ln A_i$ into the first argument of the covariance yields:

$$\text{Var}[\ln y_i] = \text{Cov}[X_i + \ln A_i, \ln y_i] = \text{Cov}[X_i, \ln y_i] + \text{Cov}[\ln A_i, \ln y_i]$$

Dividing both sides by $\text{Var}[\ln y_i]$, we immediately obtain $s_X + s_A = 1$ (where $s_X = s_k + s_h$).

4.4.1 Time-Series Decomposition for a Single Economy

Pick a specific country i and observe its trajectory over time. We start with the log-linearized production function:

$$\ln y_t = \ln A_t + \alpha \ln k_t + (1-\alpha) \ln h_t$$

Differentiating across time, we can express this relationship in terms of growth rates:

$$g_y = g_A + \alpha g_k + (1-\alpha) g_h$$

where g_x denotes the average growth rate of variable x. This additive equation allows us to directly compute the contribution share of each factor to the overall economic growth of that country:

$$s_k = \frac{\alpha g_k}{g_y}, \quad s_h = \frac{(1-\alpha) g_h}{g_y}, \quad s_A = \frac{g_A}{g_y}$$

By construction, these shares sum exactly to 1 (i.e., $s_k + s_h + s_A = 1$).

Remark (How to Compute the Average Growth from the Data?).

Take the computation of g_y as an example. A naive method is to use the endpoints: $g_y = (\ln y_T - \ln y_0)/T$. However, if the dataset is unbalanced (e.g., missing data for the exact years 0 or T), this method completely fails. Moreover, endpoint calculations are highly sensitive to business cycle fluctuations (e.g., year T happens to be a recession).

A common practice is to regress the logarithm of the variable on time t:

$$\ln y_t = \beta_0 + g_y \cdot t + \varepsilon_t$$

The OLS slope coefficient yields a robust estimate of the average growth rate g_y. Because OLS utilizes *all* available interior data points to fit the trend line, it naturally bypasses the issue of missing endpoints and smooths out short-term economic shocks.

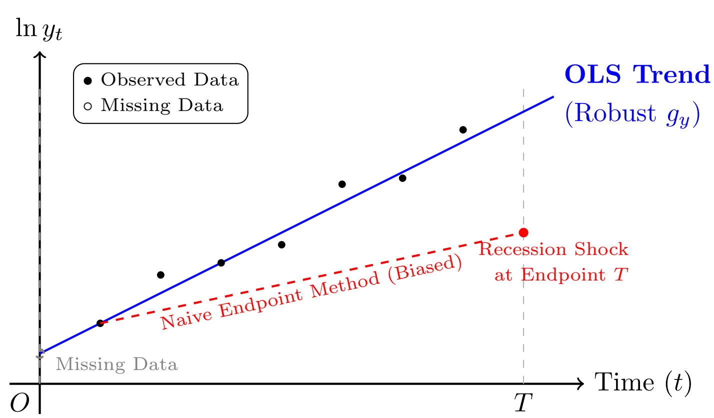

4.4.2 Cross-Sectional Decomposition of Growth Rates

The second major question in growth accounting is: **Why do some countries grow faster than others?**

To answer this, we shift our focus to the cross-sectional dispersion of growth rates.

Suppose we have computed the average growth rates for all countries, yielding the following cross-sectional dataset:

$$\{g_{y_i}, g_{k_i}, g_{h_i}\}_{i \in I}$$

Since $g_{y_i} = g_{A_i} + \alpha g_{k_i} + (1-\alpha)g_{h_i}$, we can decompose the *variance* of output growth rates ($\text{Var}[g_{y_i}]$) across countries. The methodology mirrors development accounting.

- **Method 1: The Simple Variance Ratio**
  We could naively define the contributions as:

  $$s_k^{(1)} = \frac{\text{Var}[\alpha g_{k_i}]}{\text{Var}[g_{y_i}]}, \quad s_h^{(1)} = \frac{\text{Var}[(1-\alpha)g_{h_i}]}{\text{Var}[g_{y_i}]}$$

  As discussed in Development Accounting, this method is statistically flawed. Because input growth rates and TFP growth are often correlated (e.g., fast technological adoption induces rapid capital accumulation), the covariance terms are non-zero. Consequently, the shares under Method 1 will not sum up to 100%.

- **Method 2: Covariance-Based Decomposition**
  To ensure exact additivity, the standard literature employs the covariance-based approach:

  $$s_k = \frac{\text{Cov}[\alpha g_{k_i}, g_{y_i}]}{\text{Var}[g_{y_i}]}, \quad s_h = \frac{\text{Cov}[(1-\alpha)g_{h_i}, g_{y_i}]}{\text{Var}[g_{y_i}]}, \quad s_A = \frac{\text{Cov}[g_{A_i}, g_{y_i}]}{\text{Var}[g_{y_i}]}$$

  Intuition: Once again, s_k can be visualized as the slope coefficient from an OLS regression of the capital contribution (αg_{k_i}) on the overall growth rate (g_{y_i}). This method allocates the covariance between factor accumulation and technological progress, guaranteeing that $s_k + s_h + s_A = 1$.

Remark (Chapter Summary).

- **Cobb–Douglas as the universal accounting framework.** Constant capital share α is read directly off national-accounts data, sidestepping the OLS endogeneity that would arise from estimating α via a regression of $\ln y$ on $\ln k, \ln h$.

- **Mincer wage equation operationalizes human capital.** Relative wages reflect relative productivity: $w_j/w_1 = e^{m(s_j - s_1)}$. The Mincer return m comes from labor-economics regressions; population shares come from Barro–Lee.

- **Method 2 (covariance-based) is preferred over the variance ratio.** It guarantees additivity ($s_k + s_h + s_A = 1$) by allocating the cross-term symmetrically, and admits the OLS-slope interpretation.

- **The empirical bottom line.** Across many decompositions: physical capital accounts for ∼30–40% of growth, human capital for ∼10%, and TFP for the remaining ∼50%. This motivates the Solow framework that follows: TFP, not factor accumulation, is what we ultimately need to explain.

- **Same machinery serves cross-section and time series.** Development accounting decomposes income variance across countries; growth accounting decomposes growth-rate variance over time or across countries. The mathematics is parallel.

Chapter 5

The Solow Growth Model

Remark (Notation in This Chapter).

Symbol	Meaning
A, A_t	TFP, fixed in baseline; $A_t = A_0 e^{g_A t}$ in the BGP extension
Y_t, y_t	Aggregate / per-worker output
K_t, k_t	Aggregate / per-worker capital
L	Labor force, fixed in the baseline
s	Saving rate (exogenous, constant)
δ	Depreciation rate
g_A, g_k, g_y	Steady-state growth rates of TFP, capital, output (BGP)
k^*	Steady-state per-worker capital
k_0^*	Initial capital stock placing the economy on the BGP

The Solow model is the fundamental workhorse of economic growth theory, focusing on how **capital accumulation** drives output growth over time. Its economic intuition rests on two core properties of physical capital:

- **Rivalry and Diminishing Returns:** Capital is a rival good. As we equip a fixed number of workers with more and more machines, the marginal contribution of each additional machine declines.[1]
- **Durability and Depreciation:** Capital is durable, allowing it to be accumulated over time, which generates a dynamic growth path. However, it also depreciates, meaning continuous investment is required to maintain the stock.

Key Takeaways from the Solow Model:

[1] To build intuition for why rivalry causes diminishing returns, consider the "congestion" of physical tools. If a single worker is given one unit of capital (e.g., a machine), there is no conflict. If a second unit is added, the two machines must now compete (or "conflict") for the same worker's fixed physical attention. This congestion dictates that the second machine cannot contribute as much as the first. A third machine creates even more conflict, yielding an even smaller marginal gain.

In contrast, a non-rival factor (like a software algorithm, a blueprint, or a scientific formula) does not suffer from this physical congestion. A single idea can be applied to all machines and workers simultaneously without being depleted. Therefore, an improvement in a non-rival factor acts as a direct multiplier to overall output, avoiding the trap of diminishing marginal returns.

- **Long-Run Stagnation:** Due to diminishing returns to capital, growth driven solely by capital accumulation will eventually slow down and cease.
- **Transitional Dynamics:** Economies that start with a low initial capital stock will experience strong "catch-up growth" as they converge to their steady state.

5.1 The Economic Environment

- **Time:** Discrete time, $t = 0, 1, 2, \cdots$.
- **Production Function:** In each period, the total output Y_t is produced using physical capital K_t and labor L:

$$Y_t = AF(K_t, L)$$

where the technology level A is exogenously given and fixed.

Remark.

> In reality, output depends on total *effective labor* (namely the human capital, H), which incorporates workers' education and skills (i.e., $H = h \cdot L$, where h is human capital per worker). However, in the baseline Solow model, we deliberately make two independent simplifying assumptions to isolate the mechanics of physical capital (K):
>
> 1. We abstract away from skill accumulation by assuming all workers are identical and have a constant baseline skill level ($h = 1$). Thus, human capital reduces to raw headcount: $H = L$.
> 2. We abstract away from population dynamics by assuming the raw labor force L is fixed and given (i.e., population growth rate $n = 0$).
>
> By shutting down the channels from both the education and the population, we force physical capital accumulation to be the sole driver of transitional growth in this baseline model.

- **Diminishing Returns to Capital:** The production function $F(\cdot)$ is strictly concave in capital, meaning that the marginal product of capital decreases as we accumulate more capital. Formally, we assume $F_K > 0$ and $F_{KK} < 0$.

  In addition, we assume the production function $F(\cdot)$ exhibits *Constant Returns to Scale (CRS)*. This allows us to divide both sides by L:

$$y_t = AF\left(\frac{K_t}{L}, 1\right) \equiv Af(k_t)$$

where $f'(k_t) > 0$ and $f''(k_t) < 0$, capturing the strictly positive but diminishing marginal product of capital.

- **Law of Motion for Capital**: The evolution of the aggregate capital stock is governed by:

$$K_{t+1} = (1 - \delta)K_t + I_t$$

where $\delta \in (0, 1)$ is the constant depreciation rate, and I_t is aggregate investment.

Intuitively, the future stock equals the current stock minus depreciation (outflow) plus new investment (inflow).

Remark.

> If we assume that the depreciation rate increases with the capital stock (rather than being a constant δ), it would amplify the effect of diminishing returns, making the growth path of capital even more concave and slowing down convergence further.

- **Saving Behavior:** A constant fraction of total output is saved and invested.

$$I_t = sY_t$$

where $s \in (0, 1)$ is the exogenous saving rate.

Remark (Constant Saving Behavior).

> The assumption that saving is a constant fraction of output ($I_t = sY_t$) is a purely mechanical behavioral rule, lacking strict microfoundations. It implies that households are myopic—they do not adjust their saving behavior in response to changes in the interest rate (the marginal product of capital) or expectations about future technological progress.
>
> Despite its simplicity, this assumption is still justified in macroeconomics:
>
> 1. **Empirical Validity:** It aligns with Kaldor's stylized facts, which observe that the aggregate investment-to-GDP ratio remains remarkably stable over long historical periods.
> 2. **Theoretical Robustness as a Limiting Case:** In more advanced frameworks (such as the Ramsey-Cass-Koopmans model), under some conditions, the optimal, fully forward-looking saving path actually collapses exactly to a constant saving rate.
>
> Therefore, the Solow model's assumption serves as a highly tractable, reduced-form approximation that successfully captures the core qualitative dynamics of capital accumulation without the heavy mathematical machinery of dynamic optimization.

Our goal is to characterize the dynamic path of the economy in per-worker terms:

$$\{y_t, k_t\}_{t=0}^{\infty}.$$

5.2 Solving the Model

To trace the growth path, we combine the law of motion with the saving assumption:

$$K_{t+1} = (1-\delta)K_t + sY_t$$

Dividing both sides by the constant labor force L, we obtain the fundamental difference equation of the Solow model in per-worker terms:

$$k_{t+1} = (1-\delta)k_t + sy_t = (1-\delta)k_t + sAf(k_t)$$

Because the production function $f(k_t)$ exhibits diminishing marginal returns, the transition mapping k_{t+1} as a function of k_t is strictly concave.

Given any initial capital endowment k_0, we can solve for the entire dynamic path $\{k_t\}_{t=0}^{\infty}$ by iterating on this law of motion. The intersection of this concave transition curve with the 45-degree line determines the unique steady state k^*, where the economy stops growing.

The phase diagram below visually traces out this dynamic growth path of k over time, illustrating how an economy starting from an initial $k_0 < k^*$ accumulates capital and converges to the steady state.

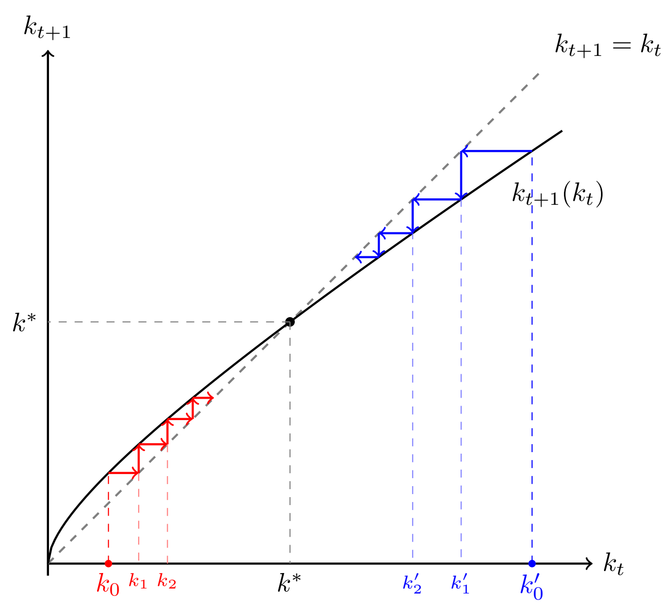

Based on the phase diagram and the mechanics of capital accumulation, we can derive two fundamental implications of the baseline Solow model:

- **No Long-Run Growth:** In the absence of continuous technological progress, the economy inevitably converges to a fixed steady state (k^*). Once there, the growth rate of capital and output per worker strictly drops to zero. This implies that capital accumulation alone cannot act as the engine of sustained long-run economic growth.

- **Catch-Up Growth (Transitional Dynamics):** Economies starting further below their steady state will grow at a faster rate. The underlying mechanism is strictly driven by the *diminishing marginal returns to capital*. When capital is scarce, its marginal product is exceptionally high, meaning investment vastly outpaces depreciation. As capital accumulates, the marginal returns diminish, causing the growth rate to gradually decelerate until the economy lands softly at the steady state.

5.3 When Variables Change

5.3.1 Technology Progress

Suppose the economy has already reached its initial steady state k^* with a technology level A. At a specific time T, technology unexpectedly and permanently improves to $A' > A$.

Because capital per worker k_t is a *state variable* (evolution is durable), it cannot jump instantly at time T. However, the technological progress causes the economy's transition curve $k_{t+1}(k_t) = (1-\delta)k_t + sAf(k_t)$ to shift up vertically. Consequently, starting from time T, the economy begins to grow rapidly. Due to the pronounced diminishing returns, this "catch-up growth" slows down quickly, and the capital stock k_t transitions smoothly yet sharply towards the new, higher steady state $k^{*\prime}$.

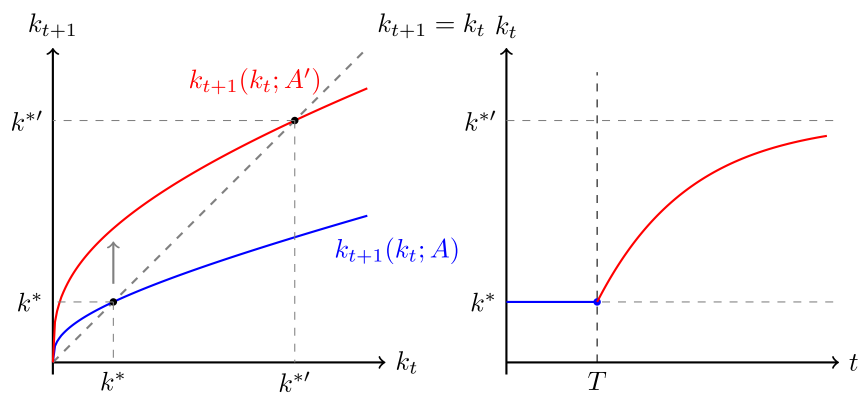

(a) Transition Diagram **(b) Time Path of Capital**

5.3.2 One-Time Population Shock

Suppose the economy has already reached its steady state k^* with a given technology level A and a constant population L. At a specific time T, the population exogenously and instantly increases to $L' > L$ (e.g., due to a sudden migration event).

Unlike a change in technology or saving rates, a one-time change in the level of population does not shift the transition curve $k_{t+1}(k_t)$. However, it directly impacts the per-worker state variable. Because the total capital stock K_T is durable and fixed at time T, the sudden increase in the denominator L causes the capital per worker k_T to jump down instantly to a lower level, $k'_T = K_T/L' < k^*$.

At this lower capital level, investment per worker now exceeds depreciation per worker. Consequently, starting from time T, the economy experiences temporarily high "catch-up growth," accumulating capital per worker until it eventually converges back to the original steady state k^*. Take a special note that the total output Y will be higher in the new steady state (due to more L), but output per worker y^* remains unchanged.

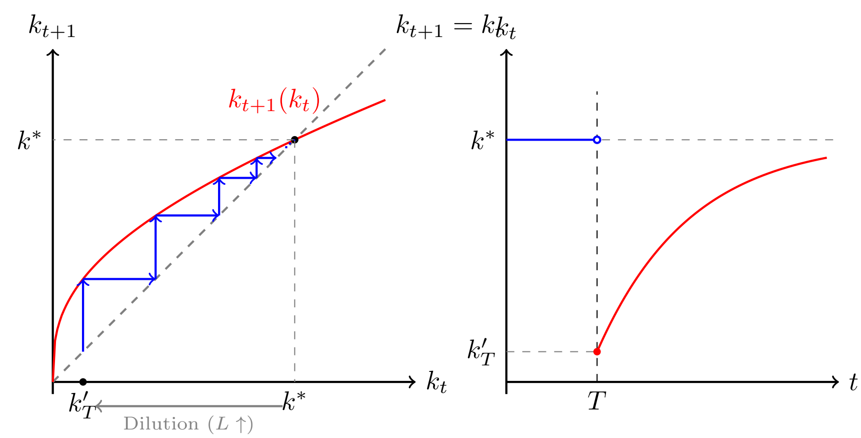

(a) Transition Diagram **(b) Time Path of Capital**

5.4 Non-Autonomous Dynamics and the Balanced Growth Path

What happens if technology grows exponentially over time, such that $A_t = A_0 e^{g_A \cdot t}$?

Now that the exogenous sequence $\{A_t\}_{t=0}^{\infty}$ is known and strictly increasing, the Solow model becomes a *non-autonomous* dynamic system.[2] We can track the sequence of capital $\{k_t\}_{t=0}^{\infty}$ by iterating the law of motion for k_{t+1}:

$$\begin{aligned} k_1 &= (1-\delta)k_0 + sA_0k_0^{\alpha} \\ k_2 &= (1-\delta)k_1 + sA_1k_1^{\alpha} \\ k_3 &= (1-\delta)k_2 + sA_2k_2^{\alpha} \\ &\vdots \end{aligned}$$

Geometrically, instead of a single curve, we plot a *family of curves* on the (k_t, k_{t+1}) phase plane:

$$k_{t+1} = (1-\delta)k_t + sA_tk_t^{\alpha}$$

Because A_t is strictly increasing with t, the transition curve shifts upward every period.

[2]In the theory of dynamical systems, a system is called *autonomous* if its law of motion does not explicitly depend on time (i.e., $x_{t+1} = f(x_t)$). The baseline Solow model without technological progress is autonomous, yielding a single, fixed transition curve. When A_t grows over time, the transition equation becomes $k_{t+1} = f(k_t, t)$. This explicit dependence on t makes the system *non-autonomous*, which graphically manifests as the transition curve shifting in every period.

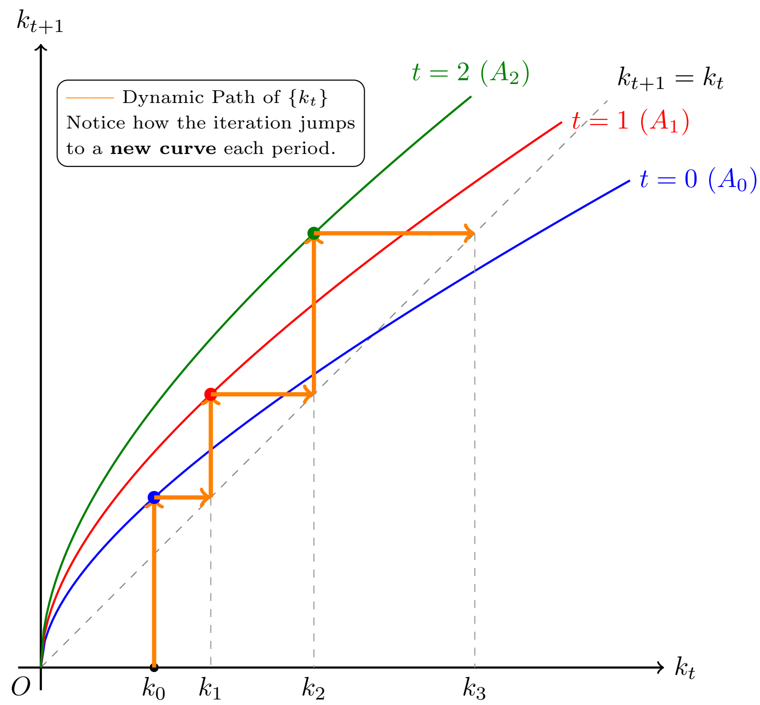

Since the transition curve is continuously shifting upward, a static steady state (where $k_{t+1} = k_t$) no longer exists. Instead, we look for a *Balanced Growth Path*: a trajectory where key macroeconomic variables (like k_t and y_t) grow at constant, time-invariant rates from the very first period.

Claim: Existence of the Balanced Growth Path

Given a constant growth rate of technology g_A, there exists a specific initial capital stock k_0^* and a constant growth rate of capital g_k, such that the sequence:

$$k_t = k_0^* e^{g_k \cdot t}$$

is an exact solution to the dynamic model for all $t \geq 0$.

Proof for Claim.

We construct the proof using the standard *Guess and Verify* method.

- **Step 1: Guess**
  Suppose the claim is true, meaning k_t grows at a constant rate g_k. We substitute $k_t = k_0 e^{g_k t}$ and $k_{t+1} = k_0 e^{g_k (t+1)}$ into the law of motion:

$$k_0 e^{g_k (t+1)} = (1-\delta) k_0 e^{g_k t} + s A_0 e^{g_A t} \left(k_0 e^{g_k t} \right)^{\alpha}$$

- **Step 2: Verify by matching coefficients**
  Divide both sides by the term $k_0 e^{g_k t}$:

$$e^{g_k} = (1-\delta) + s A_0 k_0^{\alpha-1} e^{g_A t} e^{\alpha g_k t} e^{-g_k t}$$

Rearranging the terms yields:

$$e^{g_k} + \delta - 1 = sA_0 k_0^{\alpha-1} e^{[g_A+(\alpha-1)g_k]t}$$

For this equation to hold true for *all* t (which is the definition of balanced growth), the right-hand side must be independent of time t. Therefore, the exponent on e must be exactly zero:

$$g_A + (\alpha - 1)g_k = 0 \implies g_k = \frac{g_A}{1-\alpha}$$

- **Step 3: Pin down the initial condition k_0^***
  Now that the exponential term is $e^0 = 1$, we can solve for the unique initial capital k_0^* that places the economy immediately on the BGP:

$$k_0^* = \left(\frac{sA_0}{e^{g_k} + \delta - 1}\right)^{\frac{1}{1-\alpha}}$$

As long as the economy starts at this specific k_0^*, capital will grow forever at the constant rate $g_k = g_A/(1-\alpha)$.

5.5 The Lucas Critique

Lucas Critique

Assuming structural parameters remain invariant to policy or environmental changes leads to systematic forecasting errors.

One of the most fundamental criticisms of the baseline Solow model is its assumption that the saving rate s is an exogenous constant. In reality, saving is an optimal choice made by forward-looking households. If the economic environment changes, rational agents will adjust their behavior.

Consider a positive technological shock ($A\uparrow$). In the standard Solow model, the investment curve mechanically shifts up, determining a new "naive" steady state. However, if the technological improvement alters the return to capital or lifetime income, the household's optimal saving rate s will also change.

- **Understating the Effect:** If an increase in A incentivizes households to save more ($s\uparrow$), the investment curve shifts up twice: once due to A, and again due to s. The naive Solow model will *understate* the true increase in steady-state capital.
- **Overstating the Effect:** Conversely, if higher A leads to a wealth effect that causes households to consume more and save less ($s\downarrow$), the secondary downward shift in s partially offsets the initial shock. The naive Solow model will *overstate* the true impact.

The true profundity of the Lucas Critique extends far beyond a simple complaint about the Solow model's saving rate. It sparked a methodological revolution that forced modern macroeconomics to develop rigorous *micro-foundations*.

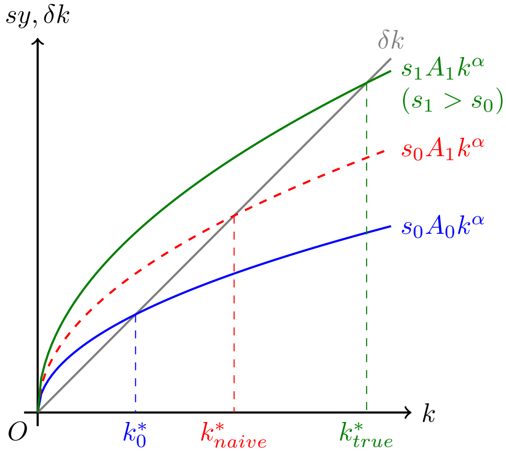

Case 1: $A \uparrow \implies s \uparrow$
Solow *understates* the effect

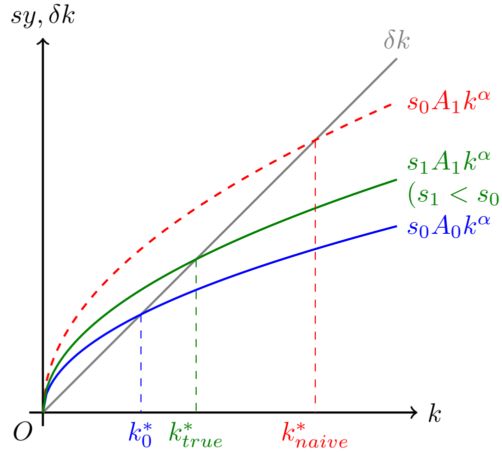

Case 2: $A \uparrow \implies s \downarrow$
Solow *overstates* the effect

Prior to Lucas (1976), macro-econometric models relied heavily on historical relationships between aggregate variables (treating variables like s or the marginal propensity to consume as structural constants). Lucas argued that these are merely *reduced-form* outcomes of underlying optimization problems. When the government changes an economic policy or when technology shifts, the constraints of the optimization problem change, and rational agents will immediately alter their behavior. Consequently, historical data becomes useless for predicting the effects of new policies.

To avoid the Lucas Critique, macroeconomic models must be built upon *deep parameters* (or structural parameters) that are genuinely invariant to policy changes. These include:

- **Preferences:** Time discount factor (β), coefficient of relative risk aversion (γ).
- **Technology:** Capital share (α), depreciation rate (δ).

This critique naturally paves the way for the next evolution in our theory: the *Neoclassical Growth Model* (often formalized as the *Ramsey-Cass-Koopmans model*). Instead of assuming an exogenous saving rule, this framework models the economy as a dynamic general equilibrium where households and firms actively optimize. In such a micro-founded setting, the saving behavior organically emerges as an endogenous outcome, perfectly immune to the Lucas Critique.

Remark (Chapter Summary).

- **No long-run growth from capital alone.** Diminishing returns to capital pin down a unique steady state k^*. Without exogenous TFP growth, capital accumulation eventually stops driving output growth.
- **Catch-up growth and conditional convergence.** An economy below its steady state grows faster than one near it. Cross-country differences in A, δ, s explain why countries converge to different k^*.
- **Balanced Growth Path.** With $A_t = A_0 e^{g_A t}$, the model admits a BGP solution where k_t, y_t grow at constant rate $g_A/(1-\alpha)$. Capital deepening still happens, but

technology is the long-run engine.

- **The Lucas Critique.** The exogenous saving rate s is a behavioral residual, not a structural parameter. Policy or technology changes that shift s are not captured in the Solow framework. This motivates the next chapter's optimization-based formulation.
- **The model as a launching pad.** The Solow setup is the basic accounting structure that all subsequent growth-and-cycles models extend, from the Ramsey–Cass–Koopmans neoclassical model (Chapter 6) through RBC and beyond.

Chapter 6

The Neoclassical Growth Model

Remark (Notation in This Chapter).

Symbol	Meaning
P_t, w_t, r_t	Final-good price, real wage, capital rental rate (often $P_t \equiv 1$)
Π_t	Profit remitted from firms to households (zero in CRS)
λ_t	Lagrange multiplier on the household's period-t budget constraint
$V(k)$	Recursive value function (per-capita formulation)
$g(k)$	Optimal policy function $k' = g(k)$
T	Bellman operator $TV(k) = \max_{k'}\{u(\cdot) + \beta V(k')\}$
k^*	Steady-state capital where $g(k^*) = k^*$
TVC	transversality condition $\lim_{t\to\infty} \beta^t u'(C_t)K_{t+1} = 0$

The fundamental departure of the Neoclassical Growth Model from the Solow model is that **growth through capital accumulation is driven by an endogenous saving rate (s)**. We no longer assume a mechanical rule; instead, we build the economy from the ground up using microeconomic optimization.

6.1 The Economic Environment

We consider an infinite-horizon discrete-time economy where time is indexed by $t = 0, 1, 2, \ldots$. The economy consists of two types of rational agents interacting in perfectly competitive markets.

- **A Representative Household (HH):**
  - Endowed with an initial capital stock K_0 and a constant labor supply L.
  - At each period t, the HH provides capital K_t and labor L to the market, earning a wage w_t and a capital rental rate r_t.
  - The HH must decide how much to consume and how much to invest/save (I_t) for the future.

- **A Representative Firm:**
  - At each period t, the firm buys (rents) capital K_t and labor L_t from the HHs.
  - The firm pays the factor prices r_t and w_t to produce output.
  - The production function is given by $AF(K_t, L_t)$, where F exhibits constant returns to scale (CRS) and has positive marginal products ($F_K > 0, F_L > 0$), and A represents the total factor productivity (TFP).
- **Market Structure**
  All markets are perfectly competitive, meaning households and firms are strict price-takers (i.e., they have zero market power and treat all prices as exogenous constants in their optimization problems). The sequence of market prices is given by:
  - Final good price: P_t[1]
  - Labor wage: w_t
  - Capital rental rate: r_t

Remark ("Representative" and "Aggregation").

> The term "representative" implicitly assumes there is a continuum of identical households and firms, both normalized to a total measure of 1. The key of the macroeconomic aggregation is *not* that the number of firms equals the number of households. Instead, it is the fundamental asymmetry between technology and preferences. Because the firm's technology is Constant Returns to Scale (scale-free), the number of firms is mathematically irrelevant. However, because the household's utility function is strictly concave (scale-dependent), we *must* normalize the continuum of households to a measure of 1. This normalization allows us to safely study a single representative individual's intertemporal choice without loss of generality.

6.2 The Household's Problem

The representative household is forward-looking and maximizes its infinite-horizon discounted lifetime utility, taking the sequence of prices $\{P_t, r_t, w_t\}_{t=0}^{\infty}$ as given.

The HH's sequence problem is formulated as:[2]

$$\max_{\{C_t, K_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u(C_t)$$

subject to the dynamic period-by-period budget constraint. Conceptually, the household's total nominal expenditure on consumption (C_t) and investment (I_t) must equal its total

[1] A "final good" refers to a product that is in its end-use stage (i.e., used directly for either household consumption C_t or capital investment I_t), as opposed to intermediate goods which are used up during the production process.

[2] Technically, the household derives utility from its individual consumption c_t. However, because we assume a representative household normalized to a measure of 1, individual consumption equals aggregate consumption ($c_t = C_t$). Standard macro notation uses capital letters directly in the utility function to reflect this equivalence.

nominal income:

$$P_t(C_t + I_t) = w_t L + r_t K_t + \Pi_t$$

Recall the fundamental law of motion for capital: tomorrow's capital is today's un-depreciated capital plus today's investment, $K_{t+1} = (1-\delta)K_t + I_t$. Substituting the implied investment $I_t = K_{t+1} - (1-\delta)K_t$ into the equation yields the standard formulation:

$$\underbrace{P_t\left[C_t + K_{t+1} - (1-\delta)K_t\right]}_{\text{Expenditure}} = \underbrace{w_t L + r_t K_t + \Pi_t}_{\text{Income}}, \quad \forall t \geq 0$$

with initial conditions and non-negativity constraints: K_0, L given, and $C_t \geq 0$ for all t.

In the HH's problem,

- $\beta \in (0, 1)$ is the time discount factor, capturing the HH's patience.
- $u(C_t)$ is the period utility function. The standard assumptions are $u' > 0$ (more is better) and $u'' < 0$ (diminishing marginal utility). The strict concavity of the utility function ($u'' < 0$) mathematically drives the desire for *endogenous consumption smoothing*.
- Π_t represents the profit remitted from the firms to the households (who are the ultimate owners of the firms).

6.3 The Firm's Problem

Unlike the household, which faces a dynamic intertemporal problem, the representative firm solves a sequence of *static* optimization problems. Capital is rented period-by-period, so the firm faces no dynamic state variables.

Taking the sequence of prices $\{P_t, w_t, r_t\}_{t=0}^{\infty}$ as given, the firm chooses input demands $\{K_t, L_t\}$ at each period t to maximize profit:

$$\Pi_t = \max_{\{K_t, L_t\}} P_t A F(K_t, L_t) - r_t K_t - w_t L_t.$$

6.4 Competitive General Equilibrium

Having specified the micro-foundations—the household's dynamic optimization and the firm's static optimization—we can now define the macroeconomic equilibrium. This definition perfectly mirrors the standard Walrasian General Equilibrium from microeconomic theory.

Definition 6.1: Competitive General Equilibrium

A *competitive general equilibrium* is defined as a sequence of prices $\{P_t, r_t, w_t\}_{t=0}^{\infty}$ and a sequence of allocations $\{C_t, L_t, K_{t+1}\}_{t=0}^{\infty}$ such that:

1. **Household Optimization:** Given the price sequence $\{P_t, r_t, w_t\}_{t=0}^{\infty}$, the allocation $\{C_t, K_{t+1}\}_{t=0}^{\infty}$ solves the representative household's utility maximization problem.

2. **Firm Optimization:** Given the price sequence $\{P_t, r_t, w_t\}_{t=0}^{\infty}$, the allocation $\{K_t, L_t\}_{t=0}^{\infty}$ solves the representative firm's profit maximization problem period by period.

3. **Markets Clear:** All markets in the economy must clear simultaneously at every period t:

   - **Goods Market Clearing:** The total supply of goods $F(K_t, L)$ equals the total demand for goods (consumption plus investment):

   $$C_t + \underbrace{K_{t+1} - (1-\delta)K_t}_{I_t} = AF(K_t, L) \quad \forall t$$

   - **Labor Market Clearing:** Firm's labor demand equals household's labor supply:

   $$L_t = L \quad \forall t$$

Remark.

The Neoclassical Growth Model is fundamentally a *real economy* model. There is no fiat money, no central bank, and no financial nominal frictions. Transactions are essentially barters (i.e., direct exchanges of goods and services without the use of money) of real goods for real factors of production.

Mathematically, if the sequences $\{P_t, w_t, r_t, C_t, L_t, K_{t+1}\}_{t=0}^{\infty}$ constitute a general equilibrium, then for any arbitrary sequence of strictly positive scalars $\{\alpha_t\}_{t=0}^{\infty} > 0$, the scaled sequence:

$$\{\alpha_t P_t, \alpha_t r_t, \alpha_t w_t, C_t, L_t, K_{t+1}\}_{t=0}^{\infty}$$

is *also* an equilibrium.

This is a direct consequence of Walrasian demand and supply functions being *homogeneous of degree zero* in prices. If both your income (wages and capital rents) and the price of consumption goods multiply by α_t, your budget constraint remains perfectly unchanged. Therefore, your optimal real allocations (C_t, K_{t+1}) do not change. Only *relative prices* (the real wage w_t/P_t and the real rental rate r_t/P_t) matter for resource allocation. This is the precise mathematical reason why many advanced macroeconomics textbooks entirely omit the price level P_t from the household's and firm's problems from the very beginning. They simply assume $P_t = 1$ for mathematical convenience, focusing exclusively on the relative real prices that actually dictate economic behavior.

As argued before, arbitrary sequence $\{\alpha_t\}$ has absolutely zero effect on real allocations (C_t, K_{t+1}). This demonstrates the *classical dichotomy*: in a frictionless real model, nominal variables (inflation) are completely divorced from real variables (growth and consumption). While the classical dichotomy holds elegantly in this frictionless theoretical model, it generally fails in the real world. In reality, the macroeconomy is rife with *nominal frictions* (e.g., sticky prices, long-term wage contracts). If prices and wages cannot adjust instantaneously and perfectly to scale with α_t, then inflation will distort relative prices, erode real purchasing power, and have profound, tangible impacts on real economic variables.

To solve for the general equilibrium analytically, we extract the FOCs from both the household's and the firm's optimization problems and combine them.

Let λ_t be the Lagrange multiplier associated with the HH's budget constraint at time t:

$$\mathcal{L} = \sum_{t=0}^{\infty} \beta^t \left\{ u(C_t) + \lambda_t \left[w_t L + (r_t + 1 - \delta)K_t + \Pi_t - C_t - K_{t+1} \right] \right\}$$

Taking derivatives with respect to the choice variables C_t and K_{t+1} yields:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial C_t} = 0 \quad &\Longrightarrow \quad u'(C_t) = \lambda_t \\ \frac{\partial \mathcal{L}}{\partial K_{t+1}} = 0 \quad &\Longrightarrow \quad \lambda_t = \beta \lambda_{t+1}(r_{t+1} + 1 - \delta) \end{aligned}$$

Combining these two first-order conditions, we obtain the Household's Euler equation:

$$u'(C_t) = \beta u'(C_{t+1})(r_{t+1} + 1 - \delta)$$

Remark (Intuition of the HH's Euler Equation).

This is the fundamental intertemporal no-arbitrage condition. At the optimum, the household must be indifferent between consuming the marginal unit today versus investing it for tomorrow.

- **LHS ($u'(C_t)$):** The marginal loss of saving. It is the utility you give up today by saving one extra unit of goods.
- **RHS ($\beta u'(C_{t+1})(r_{t+1} + 1 - \delta)$):** The marginal gain from saving, evaluated at time t. If you save one unit today, it becomes capital tomorrow, survives depreciation $(1 - \delta)$, and earns the rental rate (r_{t+1}). This total return is then converted into tomorrow's marginal utility $u'(C_{t+1})$ and discounted back to today by β.

The representative firm solves a static profit maximization problem. Taking derivatives of $\Pi_t = P_t F(K_t, L_t) - r_t K_t - w_t L_t$ with respect to L_t and K_t yields:

$$\begin{aligned} w_t &= P_t A F_L(K_t, L) \\ r_t &= P_t A F_K(K_t, L) \end{aligned}$$

Because this is a real economy where only relative prices matter, the absolute price level is purely a scalar. For analytical simplicity, it is a universal convention in computation to normalize the price of the final good to unity ($P_t = 1 \quad \forall t$). Thus, the nominal factor prices perfectly equal their real marginal products: $w_t = F_L$ and $r_t = F_K$.

We substitute the firm's real equilibrium rental rate $r_{t+1} = F_K(K_{t+1}, L)$ into the Household's Euler equation:

$$u'(C_t) = \beta u'(C_{t+1}) \left[AF_K(K_{t+1}, L) + 1 - \delta \right] \quad \cdots \text{(EE)}$$

The beauty of the Neoclassical Growth Model is its *block recursive* structure.[3] We can decouple the real allocations from the prices.

Proposition 6.2: Equilibrium Characterization

Assuming Inada conditions hold[a], a sequence of prices and allocations $\{P_t, w_t, r_t, C_t, L_t, K_{t+1}\}_{t=0}^{\infty}$ is a general equilibrium if and only if it satisfies the Transversality Condition (TVC)

$$\lim_{t\to\infty} \beta^t u'(C_t) K_{t+1} = 0,$$

along with the following two sequential blocks:

- **Block 1: Solving for Allocations (The Real Economy)**
  We completely ignore prices and solve a system of two difference equations for the two variables $\{C_t, K_{t+1}\}$:

$$\begin{cases} \text{Euler Equation:} & u'(C_t) = \beta u'(C_{t+1}) \left[AF_K(K_{t+1}, L) + 1 - \delta \right] \\ \text{Market Clearing:} & C_t + K_{t+1} - (1-\delta)K_t = AF(K_t, L) \end{cases}$$

- **Block 2: Backing out Prices**
  Once the sequence of allocations $\{C_t, K_{t+1}\}$ is found, we simply plug them into the firm's FOCs to back out the equilibrium prices $w_t = AF_L(K_t, L)$ and $r_t = AF_K(K_t, L)$ period by period.

[a] Inada conditions guarantee interior solutions $C_t > 0, K_t > 0$.

Remark (Why is the TVC Mathematically Necessary for a GE?).

> In an infinite-horizon problem, the EE only prevents short-term (one-period) deviations. The TVC ($\lim_{t\to\infty} \beta^t u'(C_t) K_{t+1} = 0$) is required to prevent the household from over-accumulating capital forever without consuming it. But a natural question arises: how does "over-accumulating capital forever" break the definition of a general equilibrium? The answer lies strictly in the household optimization condition in the GE definition.
>
> The Euler equation is only a *local* first-order condition. It merely guarantees that

[3] Mathematically, a system of equations is block recursive if it can be partitioned into blocks that can be solved sequentially. In our model, the equations determining real quantities (C_t, K_t) do not contain the price variables (w_t, r_t) once we substitute out the rental rate. Thus, we can solve for the real allocations first, and then recursively use those allocations to determine the implied prices.

the household cannot improve its lifetime utility by shifting consumption between any adjacent periods t and $t+1$. However, in an infinite-horizon setting, local optimality does not guarantee global optimality.

Suppose a proposed equilibrium path satisfies the Euler equation and market clearing, but violates the TVC, meaning $\lim_{t\to\infty} \beta^t u'(C_t)K_{t+1} > 0$. This implies that the *present utility value* of the capital stock left over at the "end of time" is strictly positive. If this is true, the household is acting sub-optimally. The household could easily construct a strictly better feasible path: simply consume a tiny bit more at $t = 0$ (which strictly increases lifetime utility) and let the entire future sequence of capital $\{K_t\}$ drift down by a small margin. Because the asymptotic present value of capital was strictly positive, this slight perpetual reduction will never cause K_t to violate the non-negativity constraint ($K_t \geq 0$) even at infinity.

Because a strictly better feasible consumption path exists, the original path *failed to maximize the household's utility.* Consequently, the household optimization condition of the GE fails. The TVC is the necessary mathematical boundary condition that rules out these "sub-optimal infinite hoarding" paths, ensuring the local Euler equation solution is indeed the global maximum.

6.5 Social Planner's Perspective

Notice that Block 1 (solving for allocations) does not contain any prices (w_t, r_t) or decentralized market features. In fact, the system of equations in Block 1 is mathematically identical to the solution of a *benevolent social planner* who wants to maximize the representative household's utility subject only to the resource constraint of the economy:

$$\max_{\{C_t, K_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t u(C_t) \quad \text{s.t.} \quad C_t + K_{t+1} = AF(K_t, L) + (1-\delta)K_t$$

Because our economic environment features no externalities, no taxes, and perfectly competitive markets, the *First Welfare Theorem* holds. The decentralized general equilibrium allocations perfectly coincide with the Pareto optimal allocations chosen by a social planner. This is why macroeconomists often skip the messy firm/household market setup entirely and just solve the social planner's problem directly to find the GE allocations.

Remark (Multiple Equilibria vs. Single Optimum).

Recall that in a standard Arrow-Debreu economy, the social planner can trace out an entire Pareto frontier of infinite optimal allocations (depending on the welfare weights assigned to different agents). A competitive equilibrium, dictated by a specific price vector and initial endowments, merely selects *one* specific point on this frontier. How can we claim they are mathematically identical here?

The resolution lies in our *representative household* assumption. Because there is effectively only *one* consolidated agent in this economy, the Pareto frontier collapses into a single, unique point. There are no distributional conflicts and no wealth transfers to consider. Consequently, the unique competitive equilibrium maps exactly 1-to-1 onto the unique Social Planner's optimum. The prices $\{w_t, r_t\}$ in this macro model do not dictate wealth distribution across heterogeneous agents; they are merely the decentralized "shadow prices" of the Planner's resource constraints.

6.5.1 Solving Planner's Problem Recursively

Let $c_t = \frac{C_t}{L}$ and $k_t = \frac{K_t}{L}$ be per-capita values.

Define the value function $V(k)$ as the maximum lifetime utility given an initial capital stock k:

$$\begin{aligned} V(k) \equiv \max_{\{c_t, k_{t+1}\}_{t=0}^{\infty}} \quad & \sum_{t=0}^{\infty} \beta^t u(c_t) \\ \text{s.t.} \quad & c_t + k_{t+1} = Af(k_t) + (1-\delta)k_t \quad \forall t \geq 0 \\ & c_t \geq 0 \quad \forall t \geq 0 \\ & k_0 = k \end{aligned}$$

We can decompose this infinite-horizon problem by separating the current period ($t = 0$) from the future ($t \geq 1$):

$$\begin{aligned} V(k) = \max_{c,k'} & \left\{ u(c) + \beta \left[\begin{aligned} \max_{\{c_t, k_{t+1}\}_{t=1}^{\infty}} \quad & \sum_{t=1}^{\infty} \beta^{t-1} u(c_t) \\ \text{s.t.} \quad & c_t + k_{t+1} = Af(k_t) + (1-\delta)k_t \quad \forall t \geq 1 \\ & c_t \geq 0 \quad \forall t \geq 1, \quad k_1 = k' \end{aligned} \right] \right\} \\ \text{s.t.} \quad & c + k' = Af(k) + (1-\delta)k \quad (t = 0) \end{aligned}$$

Recognizing that the expression inside the square brackets is exactly the definition of the value function evaluated at tomorrow's capital stock k', we obtain the recursive *Bellman equation*[4]:

$$\begin{aligned} V(k) = \max_{c,k'} \quad & u(c) + \beta V(k') \\ \text{s.t.} \quad & c + k' = Af(k) + (1-\delta)k \end{aligned}$$

From this we define the policy function as $k'^* = g(k)$, which is the solution (i.e., the maximizer) to the Bellman equation.

[4]In dynamic programming, the Bellman equation (or dynamic programming equation) is a functional equation based on Bellman's *Principle of Optimality.* It states that an optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision. Mathematically, it translates an infinite-dimensional sequence problem into a two-period recursive problem: maximizing the sum of the current flow payoff and the discounted continuation value evaluated at the optimally chosen next-period state.

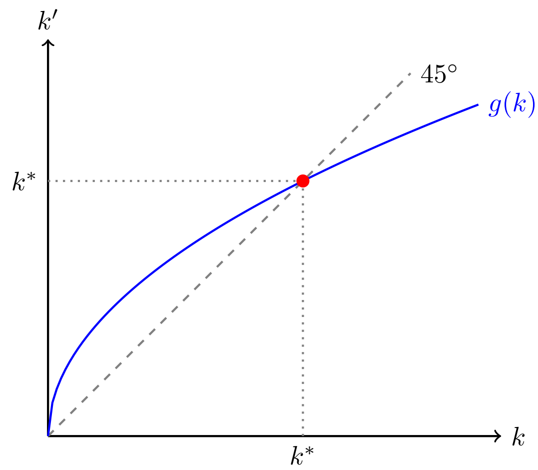

This phase diagram illustrates the dynamic evolution of capital via the policy function $k' = g(k)$. The horizontal axis represents the current capital stock k, and the vertical axis represents tomorrow's capital stock k'. The dashed 45° line represents the locus of points where $k' = k$ (i.e., capital remains constant). The intersection of the policy function $g(k)$ and the 45° line strictly determines the *steady state* of the economy, denoted as k^*. Because the policy function $g(k)$ is strictly concave and crosses the 45° line from above, the steady state is globally stable: for any initial $k_0 > 0$, the economy will dynamically converge to k^*.

Remark (The (c, k) Phase Diagram and the Saddle Path).

A complementary view of the same dynamics plots the joint evolution of consumption and capital in the (k, c) plane. Two loci organize the picture:

- The $\Delta c = 0$ **locus** (Euler-equation steady state): the locus of states at which marginal utility is unchanged from one period to the next. From $u'(c_t) = \beta\, u'(c_{t+1})\,[Af'(k_{t+1}) + 1 - \delta]$, this requires $Af'(k) = (1/\beta) - 1 + \delta$, which pins down a single capital stock $k = k^*$ *independent of* c. The locus is therefore a *vertical line*.

- The $\Delta k = 0$ **locus** (resource-constraint steady state): the locus where investment exactly replaces depreciation. From the resource constraint $c + k' = Af(k) + (1 - \delta)k$, setting $k' = k$ gives $c = Af(k) - \delta k$. This locus is a *hump-shaped curve* in (k, c) space, peaking at the Golden-Rule capital level.

The two loci cross at the unique steady state (k^*, c^*). Around this point, the dynamics organize into four regions, each with a characteristic direction of motion. Crucially, only *one* trajectory through any initial k_0 converges to the steady state—the so-called **saddle path**. All other trajectories diverge: capital either accumulates without bound (consumption too low, asymptotically violating the transversality condition by leaving wealth on the table) or runs to zero (consumption too high, violating non-negativity).

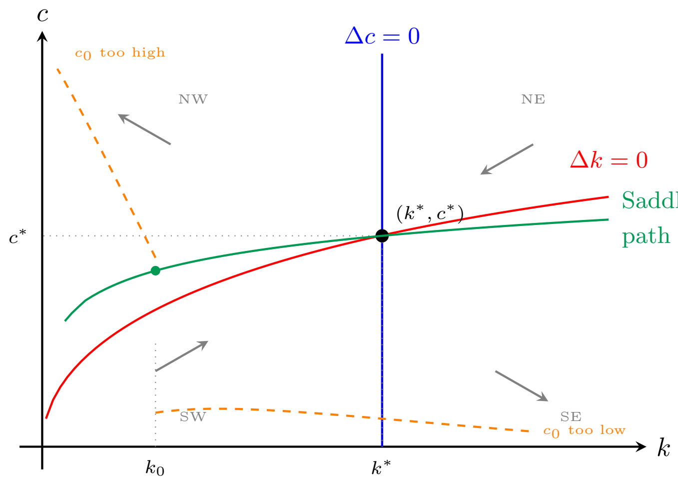

Three features of the picture are worth flagging.

- **Saddle stability is a generic feature of optimal-control problems.** The Euler equation is forward-looking (consumption today depends on consumption tomorrow), the capital constraint is backward-looking (capital tomorrow depends on capital today). The two loci cross transversally at (k^*, c^*); one eigenvalue is stable, one is unstable. Hence a one-dimensional stable manifold—the saddle path.

- **The saddle path is the global solution.** Given k_0, the optimal initial consumption c_0 is the unique value such that (k_0, c_0) lies on the saddle path. Any other choice eventually violates either the resource constraint (k goes to zero) or the transversality condition (k grows without bound while leaving utility on the table).

- **The recursive solution and the saddle path describe the same trajectory.** The discrete-time policy function $k' = g(k)$ from the Bellman equation is exactly the projection of the saddle path onto the (k, k') plane: starting from k_0, applying g repeatedly traces out the same sequence $\{k_t\}$ as following the saddle path in the (k, c) plane.

6.5.2 Solving the Model Numerically: Value Function Iteration

The mathematical justification for Value Function Iteration relies on functional analysis, specifically the *Banach Fixed-Point Theorem*. Let T be the Bellman operator mapping a value function V into a new value function TV:

$$(TV)(k) = \max_{k'} \{u(Af(k) + (1 - \delta)k - k') + \beta V(k')\}$$

According to *Blackwell's Sufficient Conditions*, because the operator T satisfies *monotonicity* (if $V \geq W$, then $TV \geq TW$) and *discounting* (for any constant $a \geq 0$, $T(V + a) \leq TV + \beta a$), the Bellman operator T is a **contraction mapping** with modulus $\beta \in (0, 1)$.

Theorem 6.3: Banach Fixed Point Theorem

- *Existence and Uniqueness:* There exists a unique true value function V^* such that $TV^* = V^*$ (a unique fixed point).
- *Global Convergence:* For *any* initial guess $V^{(0)}$, the sequence generated by $V^{(i+1)} = TV^{(i)}$ will converge to the true value function V^* uniformly as $i \to \infty$. This is the theoretical bedrock that allows us to safely initialize VFI with an arbitrary guess like $V^{(0)}(k) = 0$.

To solve the recursive Bellman equation computationally, we use an algorithm known as *Value Function Iteration* (VFI). Since a computer cannot process an infinite-dimensional continuous function, we must approach the problem through discretization and recursive iteration.

The standard VFI algorithm proceeds as follows:

Algorithm: Value Function Iteration

1. **Set a grid on the state space:**
   Choose a finite set of n discrete grid points for the capital stock:
   $$K = \{k_1, k_2, \ldots, k_n\}$$
2. **Start with an initial guess $V^{(0)}$:**
   Initialize the value function for all points on the grid. A common and mathematically valid starting point is zero:
   $$V^{(0)}(k) = 0, \quad \forall k \in K$$
3. **Update the value function using the Bellman Equation:**
   For each $k \in K$, compute the updated value function $V^{(i+1)}(k)$ by solving the maximization problem on the right-hand side (RHS). In the most basic version of the algorithm, we restrict the choice variable k' to also be chosen from the discrete grid K:
   $$V^{(i+1)}(k) = \max_{k' \in K} \left\{ u \underbrace{\left(Af(k) + (1-\delta)k - k'\right)}_{\text{consumption } c} + \beta V^{(i)}(k') \right\}$$
   Note: The underbrace does not imply c is a constant. It simply indicates that the term inside the parenthesis is the household's consumption level, substituting out the resource constraint.
4. **Check for convergence:**
   Evaluate the maximum absolute difference (the sup-norm) between the new

and old value functions over the entire grid:

$$\max_{k \in K} \left| V^{(i+1)}(k) - V^{(i)}(k) \right| < \epsilon$$

where ϵ is a pre-specified small tolerance level (e.g., 10^{-6}).[a]

- If the condition holds, the value function has converged. Stop the iteration.
- Otherwise, set $i \leftarrow i + 1$ and return to Step 3.

5. **Extract the Policy Function**

   Once the value function has converged to V^*, finding the optimal policy function $k' = g(k)$ requires just one final pass through the state space. For each grid point $k \in K$, we find the k' that maximizes the RHS of the Bellman equation using the converged V^*:

   $$g(k) = \arg\max_{k' \in K} \{u(Af(k) + (1 - \delta)k - k') + \beta V^*(k')\}$$

   In coding terms, instead of storing the `max` value as we did during the iteration, we now record the `argmax` (the optimizer) to recover the optimal path of capital.

[a]**Coding Practice:** In programming languages like R, Python, or MATLAB, this step is typically implemented using a `while` loop. To prevent an infinite loop if the algorithm fails to converge (e.g., due to poor calibration or grid issues), you should always code a fail-safe maximum number of iterations: `while (error > tol) && (iter < max_iter)`. If the loop terminates because it hits `max_iter`, the code should print a warning.

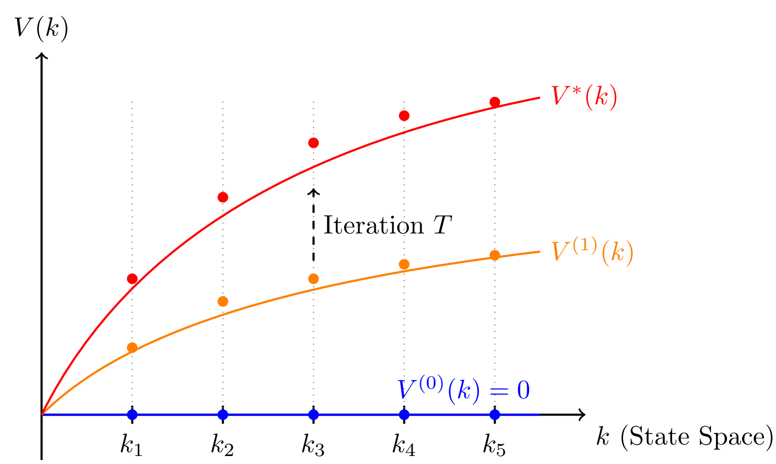

Remark (Continuous Choice, Interpolation, and the Curse of Dimensionality).

In the basic algorithm above, both the current state k and the choice variable k' are restricted to the discrete grid K. However, this purely discrete approach creates a fundamental problem: restricting k' to a coarse grid yields highly inaccurate policy functions.

To improve precision without increasing the grid size, advanced numerical solvers allow the choice variable k' to be *continuous*, even though the state space k remains discrete. When a continuous optimizer searches for the optimal k', it will likely select a value that falls *between* two of our pre-defined grid points. Since the computer only stores the value function $V^{(i)}$ at exact grid points, it must guess the value of $V^{(i)}(k')$ using **interpolation**. Common interpolation methods include:

- **Step function (Nearest-neighbor):** The worst possible method. It rounds k' to the nearest grid point, creating artificial "flat spots" (zero derivatives) and discontinuities. This destroys the strict concavity of the value function, rendering derivative-based First-Order Conditions (FOCs) entirely useless.
- **Piecewise linear interpolation:** Simple and robust, but introduces kinks (non-differentiability) at the grid points.
- **Spline interpolation** (e.g., cubic splines): The ideal method. It preserves the smoothness and strict concavity of the value function, allowing for fast, derivative-based optimization.[a]

Then a natural question arises: *Why not simply use a massive discrete grid and avoid interpolation altogether?*
Doing so would trigger the *Curse of Dimensionality.* As the number of state variables in a macroeconomic model increases (e.g., adding employment status, idiosyncratic productivity shocks, etc.), the total number of required grid points grows exponentially. For instance, a model with 3 state variables each having 100 grid points requires evaluating $100^3 = 1,000,000$ points per iteration. This exponential explosion in memory usage and computational time forces macroeconomists to keep grid sizes small and rely heavily on continuous choice and sophisticated interpolation methods.

[a]**Spline Interpolation:** A spline is a mathematical function defined piecewise by polynomials. A *cubic spline* connects the predefined grid points (knots) using third-degree polynomials, matching not just the function values, but also the first and second derivatives at each knot. This ensures the resulting approximated function is globally continuous and smooth.

Remark (Grid Design: Anchoring the Grid Around the Steady State).

A question that is logically prior to the choice of interpolation method is: *where should the grid points be placed?* The answer is not arbitrary—it depends on both the structure of the model and the specific question being asked.

Step 1: Find the steady state first. Before constructing the grid, it is useful to characterize the steady state k^* at which $g(k^*) = k^*$. Depending on the model, this may be available analytically (e.g., from the Euler equation at steady state) or may itself require a numerical root-finding step. Either way, k^* serves as the natural anchor for grid design.

Step 2: Align the grid with the research question.

- *Transition dynamics.* If the goal is to study how an economy converges from a low initial capital stock $k_0 \ll k^*$ toward the steady state—a common question in development economics—then the grid should span $[k_{\min}, k^*]$ (or slightly beyond k^* for symmetry), with points concentrated in the range where the value function and policy function have the most curvature, typically near $k_{\min}$.

- *Fluctuations around the steady state.* If instead the goal is to study business-cycle dynamics—small perturbations around k^* induced by productivity shocks—then a narrower grid centered tightly around k^* is more appropriate. Placing many grid points far from k^* wastes computational resources on regions the economy will never visit in the relevant simulation.

*A uniformly spaced grid allocates resolution evenly, but the value function and policy function are often much more curved at low k (where the Inada condition implies $f'(k) \to \infty$) and nearly linear near the steady state. A **log-spaced grid**—evenly spaced in $\log k$ rather than k—allocates more points to the low-k region where precision matters most, without requiring an excessively large total number of grid points.*

Having developed the neoclassical theory of growth—both its analytical machinery and its numerical implementation—we now turn it on the data. The next chapter asks whether the model's signature prediction of convergence actually holds across countries, and what modifications can reconcile theory with the cross-country evidence.

Remark (Chapter Summary).

- **Endogenous saving via household optimization.** The neoclassical model replaces Solow's exogenous saving rate with a forward-looking household solving an infinite-horizon dynamic problem. The Euler equation pins down the consumption path.

- **Block-recursive structure.** Real allocations (C, K) can be solved separately from prices (w, r), because the equilibrium is decoupled by the firm's static FOCs. This is the formal underpinning for using the social planner's problem as a shortcut.

- **First Welfare Theorem in action.** With complete markets, no externalities, and a representative household, the competitive equilibrium coincides with the social planner's optimum. Solving the planner's problem is therefore valid.

- **Recursive formulation and VFI.** The Bellman equation, paired with Banach's Fixed Point Theorem, guarantees a unique value function reached by Value Function Iteration from any initial guess. Implementation requires a state-space grid and a choice between discrete and continuous control.

- **The saddle path.** In the (c, k) phase diagram, the unique trajectory through any k_0 that converges to the steady state is the saddle path; all other trajectories violate either non-negativity or the transversality condition. This is the canonical picture every macroeconomist carries around.

- **Bridge to data and to cycles.** Chapter 7 confronts this model with the cross-country evidence; Chapters 8–9 add stochastic shocks to convert it into a business-cycle model.

Chapter 7

Neoclassical Growth vs. Data

Remark (Notation in This Chapter).

Symbol	Meaning
$g_y^{1960\text{-}1990}$	Average growth rate of y over the indicated cross-country sample window
k_L^*, k_H^*	Low and high stable steady states (Mechanism I: non-convexities)
$\hat{k}$	Unstable threshold / tipping point separating their basins of attraction
A_L, A_H	Low and high TFP levels (Mechanism II: structural heterogeneity)

The preceding chapters developed the neoclassical growth framework—from the Solow model to its recursive formulation in the Ramsey–Cass–Koopmans model, together with computational tools such as Value Function Iteration. We now confront the theory with cross-country data. The central question is: does the model's signature prediction of convergence actually hold? If not, what modifications can reconcile theory with evidence?

7.1 The Convergence Puzzle

A central prediction of the Solow model is *unconditional convergence*: because the marginal product of capital is high when capital is scarce, poorer countries should grow faster than richer ones and eventually catch up. Formally, under globally diminishing returns, the capital accumulation equation

$$k_{t+1} = sf(k_t) + (1 - \delta)k_t$$

is a strictly concave function of k_t, crossing the 45-degree line exactly once at a unique, globally stable steady state k^*. Regardless of the initial capital stock $k_0 > 0$, the economy eventually converges to k^*.

The cross-country data, however, paint a considerably more complicated picture. The diagrams below illustrate, in stylized form, the joint distribution of initial income and subsequent growth rates across two historical periods.

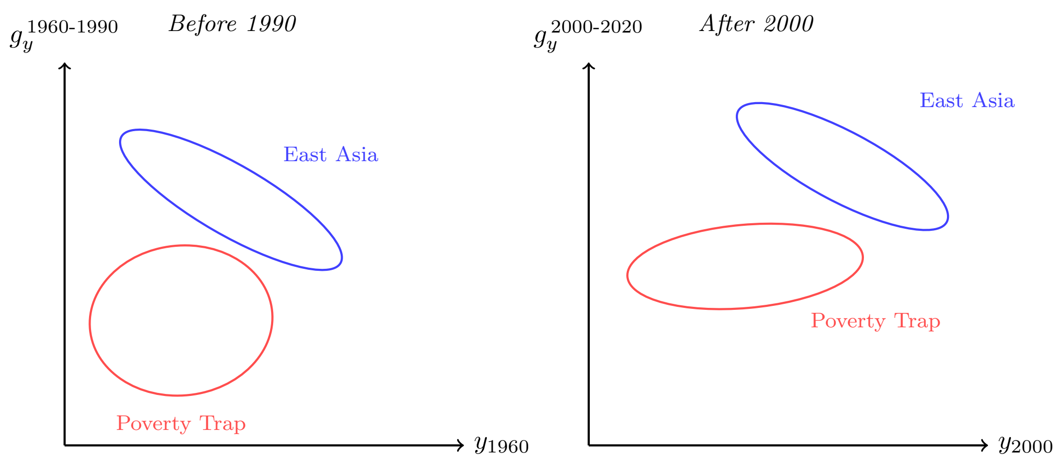

Two features of the data stand out. First, prior to 1990, the cross-country distribution of (initial income, growth rate) pairs is strikingly *bimodal*. A cluster of East Asian economies—South Korea, Taiwan, Singapore, Hong Kong—achieved rapid catch-up growth of precisely the kind the Solow model predicts, growing fast from a moderate initial base. Yet a much larger cluster of low-income economies, concentrated in sub-Saharan Africa and parts of South Asia, exhibited persistently low or near-zero growth rates. These economies did not converge; they appeared stuck.

Second, and more troublingly, the bifurcation persists after 2000. The economies that were poor in 1960 and failed to grow remained poor in 2000, and their growth performance continued to lag. This pattern is what we call a *poverty trap*: an apparently self-reinforcing state in which low income begets low investment, which begets low income.

The poverty trap constitutes a direct challenge to the basic neoclassical framework. Under globally diminishing returns, there is no reason for any economy with positive capital to stagnate indefinitely; every capital-scarce economy should enjoy high marginal returns and grow. To rationalize the data, we need to enrich the theory. We consider two distinct theoretical mechanisms.

7.2 Two Mechanisms for Poverty Traps

7.2.1 Mechanism I: Non-Convexities and Multiple Steady States

The first approach preserves the neoclassical structure but relaxes the assumption of globally diminishing returns. If returns to capital are *locally increasing* over some intermediate range—due, for example, to fixed costs of technology adoption, learning-by-doing externalities, or coordination failures across sectors—the law of motion for capital becomes *S-shaped* rather than globally concave.

The intuition is as follows. At very low capital levels, the economy cannot generate enough surplus to invest in the technologies and complementary inputs needed for growth: returns are low, investment is low, and the economy stagnates. But once capital crosses a critical threshold, positive externalities and complementarities kick in, and growth accelerates. Eventually, at high capital levels, the usual diminishing returns reassert themselves, and the economy settles into a high steady state.

Formally, the resulting law of motion $k_{t+1} = g(k_t)$ crosses the 45-degree line *three times*, yielding multiple steady states:

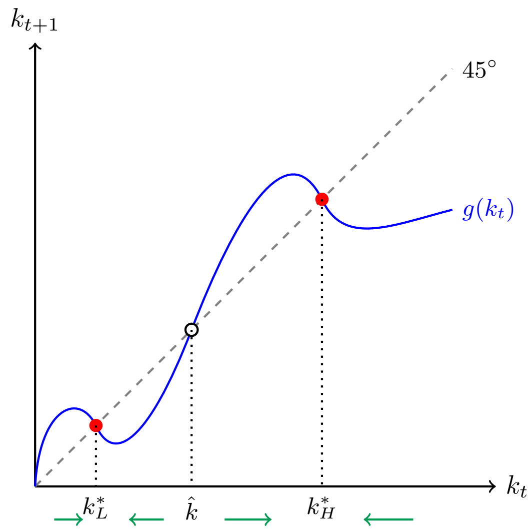

The three intersections yield two stable and one unstable steady state:

- k_L^*: the **low stable steady state**—the poverty trap. For any initial capital $k_0 < \hat{k}$, the economy converges to k_L^*. The low steady state is stable because, on both sides of k_L^*, the dynamics push the economy back: for $k < k_L^*$, $g(k) > k$ so capital rises; for $k_L^* < k < \hat{k}$, $g(k) < k$ so capital falls.
- $\hat{k}$: an **unstable steady state**—the tipping point. It is unstable because $g'(\hat{k}) > 1$: any small perturbation sends the economy either toward k_L^* (if $k_0 < \hat{k}$) or toward k_H^* (if $k_0 > \hat{k}$). Stability requires $g'(k^*) < 1$ at a steady state, which fails here.
- k_H^*: the **high stable steady state**—the developed equilibrium. For any initial capital $k_0 > \hat{k}$, the economy converges to k_H^*.

On this view, the East Asian miracle was the story of economies that—through a combination of historical circumstances, early industrial policy, and perhaps favorable initial conditions—found themselves above the tipping point $\hat{k}$ and subsequently converged to k_H^*. The stagnating economies of sub-Saharan Africa, by contrast, remained in the basin of attraction of k_L^*.

7.2.2 Mechanism II: Structural Differences Across Economies

The second mechanism takes a different route entirely. Rather than introducing non-convexities, it maintains globally diminishing returns but allows for **heterogeneity in total factor productivity (TFP)** across countries. Countries differ in institutions, geography, human capital, governance, and technological access—all of which shift the productivity parameter A in the production function and, consequently, shift the entire law of motion.

Formally, suppose two economies share the same functional form f but differ in TFP:

$$k_{t+1} = sA_i f(k_t) + (1-\delta)k_t, \qquad i \in \{H, L\}, \quad A_H > A_L.$$

Each economy has a *unique*, globally stable steady state determined by where its own law of motion intersects the 45-degree line. Because $A_H > A_L$, economy H's law of motion lies strictly above economy L's, yielding $k_H^* > k_L^*$.

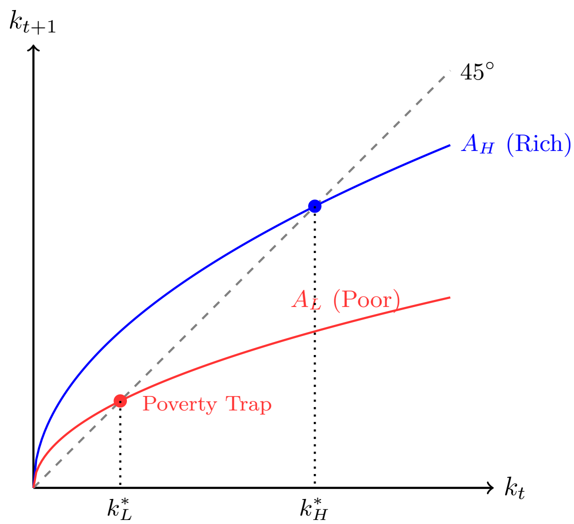

Under this mechanism, the poverty trap is not a trap in the traditional sense of a basin of attraction with a tipping point. The poor economy is simply converging to a *different*, structurally lower steady state. No matter how long we wait, it will not catch up to the rich economy—unless its structural fundamentals (institutions, technology, human capital) improve.

The key empirical implication is **conditional convergence**: controlling for the determinants of the steady state, economies with lower capital stocks grow faster. This is entirely consistent with the neoclassical model, and indeed the empirical literature on cross-country growth regressions finds considerably stronger support for conditional convergence than for the unconditional variety.

The two mechanisms are conceptually distinct but not mutually exclusive. A country may face both non-convexities and structural disadvantages. Disentangling their relative importance is, as we shall see, an extraordinarily difficult empirical task.

7.3 Empirical Challenges: The Big Push

7.3.1 Concept and Motivation

A natural policy response to the poverty trap—particularly if one believes Mechanism I—is the **Big Push**: a large-scale, coordinated injection of capital into a low-income economy with the aim of pushing k_0 above the tipping point $\hat{k}$ and into the basin of attraction of k_H^*. The idea was formalized by Murphy et al. (1989) in a model of industrialization with increasing returns to scale, building on Rosenstein-Rodan's (1943) earlier insight that coordinated, large-scale investment could overcome the coordination failures preventing take-off.

The policy logic is compelling under Mechanism I: the required intervention may be large, but it need not be permanent. Once k_0 is pushed above $\hat{k}$, the economy accumulates on its own toward k_H^*, and the intervention can be withdrawn.

7.3.2 The Identification Problem

The empirical challenge is severe: *regardless of whether a Big Push succeeds or fails, the outcome is consistent with both mechanisms.* This creates a fundamental identification problem.

Case 1: The push fails. The economy receives a large capital infusion but eventually reverts to low income.

- Under Mechanism I: the injection may have been substantial but still insufficient to push k_0 past $\hat{k}$. We do not know where $\hat{k}$ lies, and a failed push cannot tell us the economy received *enough* capital to test whether the threshold exists.
- Under Mechanism II: capital alone cannot raise the steady state if TFP remains unchanged. The economy was always converging to a structurally low k_L^*, and the capital simply depreciated.

A failed push does not allow us to distinguish between the two explanations.

Case 2: The push succeeds. The economy receives a large capital infusion and subsequently exhibits rapid, sustained growth.

- Under Mechanism I: this is the predicted outcome—the economy crossed $\hat{k}$ and is now converging to k_H^*.
- Under Mechanism II: large-scale interventions rarely deliver capital alone. They typically come packaged with technology transfer, institutional reform, managerial expertise, and infrastructure investment—all of which raise A_L toward A_H, shifting the entire law of motion upward. Under Mechanism II, the economy converges to a newly elevated k^*; the capital infusion was the vehicle, but the structural improvement was the cause.

A successful push does not cleanly identify Mechanism I as the correct explanation either.

This identification problem is deep and not merely a matter of better data. The two mechanisms are, in many empirically relevant situations, *observationally equivalent*: the same cross-country patterns of growth are consistent with either a world of multiple stable steady states or a world of heterogeneous but unique steady states. This is one of the central methodological difficulties in empirical development economics.

7.3.3 Policy Implications of the Two Mechanisms

Despite the identification challenge, the two mechanisms carry sharply different policy prescriptions.

Under Mechanism I, the key lever is *crossing the threshold.* A targeted, temporary capital injection—if sufficiently large—can permanently transform a poor economy's trajectory. The intervention must be big, but it need not be sustained indefinitely; the economy's own internal dynamics carry it to k_H^* once it is above $\hat{k}$.

Under Mechanism II, a capital injection alone is insufficient. The structural sources of low TFP—weak institutions, geographic barriers, low human capital, or technological backwardness—must be addressed directly. Capital without structural reform will simply depreciate, and the economy will return to k_L^*. Policy must target the fundamentals, not just the capital stock.

A plausible view is that both mechanisms operate simultaneously: poverty traps partly reflect non-convexities and partly reflect structural disadvantage, with the relative importance varying by country and historical context. This makes the design of effective development policy inherently difficult and context-dependent.

Remark (Development Economics and the Experimental Approach).

A strand of the development economics literature, closely associated with Esther Duflo and Abhijit Banerjee (Nobel Prize 2019), has attempted to make progress on these questions through **randomized controlled trials (RCTs)**. Rather than studying Big Push experiments at the national level—where identification is nearly impossible given the confounding of capital, technology, and institutions—this approach evaluates targeted interventions at the household or village level, aiming to isolate the causal effect of capital infusions on income and investment behavior.

While greatly improving causal identification within their specific context, micro-level RCTs face their own important limitations. Interventions at small scales may fail to replicate the *general equilibrium* complementarities, coordination effects, and agglomeration economies that lie at the heart of the Big Push mechanism. A village-level cash transfer is simply not the same as a national industrialization drive. The external validity of micro-level experiments for macro-level policy questions therefore remains an active and contentious area of debate.

This chapter connects directly to the broader convergence debate. If Mechanism II is correct, convergence is conditional on structural fundamentals: economies with similar TFP levels converge to similar steady states, consistent with the neoclassical model and the empirical finding of conditional convergence. If Mechanism I is correct, convergence depends on initial conditions in a fundamentally more path-dependent way: history matters, and similar structural fundamentals may still give rise to different long-run outcomes depending on where the economy started relative to $\hat{k}$.

Remark (Chapter Summary).

- **The convergence puzzle.** Cross-country data on (initial income, growth rate) are bimodal: a high-growth East Asian cluster and a stagnant cluster of low-income economies that fail to converge. The basic neoclassical model cannot rationalize the latter under globally diminishing returns.

- **Two competing mechanisms.** (I) Non-convexities producing an S-shaped law of motion with multiple stable steady states $\{k_L^*, k_H^*\}$ and an unstable threshold $\hat{k}$. (II) TFP heterogeneity: different countries converge to structurally different unique steady states.

- **Identification problem.** Under either mechanism, the same Big Push intervention can succeed or fail. The two mechanisms are observationally equivalent in many empirically relevant settings, which makes development policy genuinely difficult.

- **Conditional vs. unconditional convergence.** Mechanism II implies conditional convergence (controlling for fundamentals). Mechanism I implies path dependence: history matters, and similar fundamentals can produce different long-run outcomes depending on initial conditions relative to $\hat{k}$.
- **Modern empirics with RCTs.** The Banerjee–Duflo experimental approach addresses identification at the household or village level, but the external validity for macro-level Big Push interventions remains contested.

Chapter 8

Real Business Cycles

Remark (Notation in This Chapter).

Symbol	Meaning
y_t^T, y_t^C	Trend and cyclical components of $y_t \equiv \ln(Y_t/\text{pop}_t)$
$a_t = \ln A_t$	Log-TFP
ρ	Persistence of log-TFP (AR(1))
σ_ε	Standard deviation of TFP innovations
Π, π_{ij}	Markov transition matrix and its entries (discrete-state version)
$V(a, k)$	Planner's value function
$g(a, k)$	Optimal policy $k' = g(a, k)$
$\text{CV}(x_t)$	coefficient of variation of x_t, used as a unit-free volatility measure
B	Disutility-of-labor scale parameter (endogenous-labor extension)
ν	Frisch elasticity of labor supply
MPL, MPK	Marginal products of labor and capital, $e^a F_L(K, L)$ and $e^a F_K(K, L)$
EE_K, EE_L	Inter- and intra-temporal Euler equations (endogenous-labor extension)

The preceding chapters focused on the **long-run** behavior of the economy: what determines the level and growth rate of output over decades. We now shift our attention to **short-run fluctuations**—the business cycle. The central question changes from "Why do some countries grow faster than others?" to "Why does output fluctuate around its long-run trend, and what accounts for the patterns we observe in those fluctuations?"

8.1 Defining the Business Cycle

To study business cycles, we must first separate the data into what is "trend" and what is "cycle." Consider the time series of log real GDP per capita, $y_t \equiv \ln(Y_t/\text{pop}_t)$. The raw data exhibits a clear upward trend (driven by long-run growth) with fluctuations around it. We decompose:

$$y_t = y_t^T + y_t^C,$$

where y_t^T is the **trend component** and $y_t^C \equiv y_t - y_t^T$ is the **cyclical component** (the percentage deviation of output from trend).

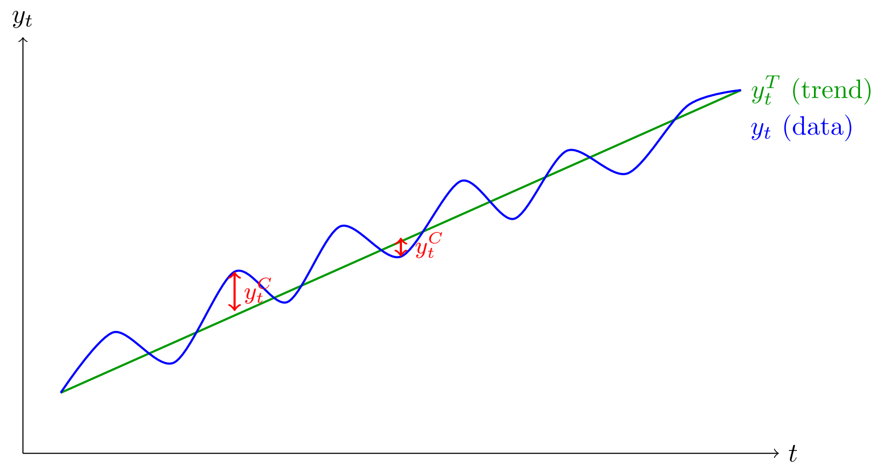

Remark (The Trend-Cycle Decomposition is Not Unique).

The decomposition above is only as meaningful as the method used to extract the trend. Different detrending procedures—such as fitting a log-linear regression, the Hodrick-Prescott (HP) filter, or the Baxter-King band-pass filter—will produce different trend series $\{y_t^T\}_{t=0}^T$ and therefore different cyclical components $\{y_t^C\}_{t=0}^T$. The "business cycle" is not a directly observable object; it is defined by how we choose to remove the trend. This is an important caveat: empirical facts about business cycles are always conditional on the detrending method.

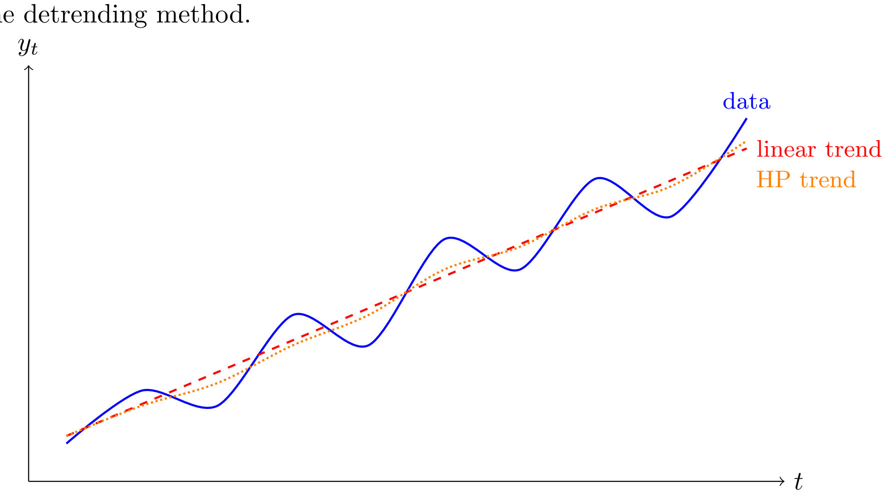

A **recession** is informally defined as two or more consecutive quarters in which *real GDP falls relative to the previous quarter* (i.e., negative quarter-over-quarter growth), or equivalently, a sustained rise in the unemployment rate. This is a statement about the *level* of GDP, not about its deviation from trend. In the United States, the official arbiter is the NBER Business Cycle Dating Committee, which considers a broader set of indicators—with particular weight on employment—rather than mechanically applying the two-quarter rule. Postwar U.S. recessions have averaged roughly 11 months in length, with a typical peak-to-trough decline of about 3% in real GDP and a 2 percentage point increase in the unemployment rate.

Remark (Why "Real"?).

> The qualifier "real" in *Real Business Cycle* contrasts the theory with earlier *monetary* business cycle theories (e.g., Friedman–Schwartz, Lucas's island model) that traced fluctuations to monetary surprises. RBC instead generates cycles from **real** (i.e., non-nominal) shocks—shocks to technology, preferences, or government spending—propagated through real economic mechanisms (capital accumulation, intertemporal substitution, labor supply). In its purest form, the RBC model has no money, no nominal rigidities, and no role for monetary policy. Cycles are not pathologies to be smoothed away; they are the optimal response of a frictionless economy to exogenous productivity changes.

8.2 Three Key Facts about Business Cycles

Any model that aspires to explain business cycles must be consistent with the following empirical regularities, documented extensively by King and Rebelo (1999). Underlying all three is a more basic "Fact 0": **the major macroeconomic aggregates—consumption, investment, hours, productivity—all comove positively with output**. A successful business cycle model must reproduce this comovement, not just match individual volatilities.

Fact 8.1: Investment is more cyclical than output

The cyclical component of investment has a much larger variance than that of output. More generally, the more durable a good is, the more cyclical its expenditure: housing and cars fluctuate far more than clothing or services. Nondurable consumption, by contrast, is *less* cyclical than output (Bils et al., 2012).

Fact 8.2: Productivity (the Solow Residual) is pro-cyclical

Total factor productivity, measured as the Solow residual, comoves positively with output over the cycle. Output tends to be high when measured productivity is high, and vice versa.

Remark (The Solow Residual).

> Given a Cobb-Douglas production function $Y_t = A_t K_t^{\alpha} L_t^{1-\alpha}$, the **Solow residual** is the component of output growth that cannot be accounted for by the growth of measured inputs:
>
> $$\ln A_t = \ln Y_t - \alpha \ln K_t - (1-\alpha)\ln L_t.$$
>
> It is often interpreted as a measure of total factor productivity (TFP). In practice, α is calibrated to the capital share of income (roughly 1/3 in the U.S.), and K_t, L_t, Y_t are taken from national accounts data. The residual then captures everything that affects output beyond capital and labor—including technology, institutions, measurement error, and capacity utilization.

Fact 8.3: Total hours worked are about as cyclical as output

Aggregate labor input fluctuates with roughly the same amplitude as output. Approximately 2/3 of the variation in total hours comes from the *extensive margin* (the number of employed workers), and 1/3 from the *intensive margin* (hours per worker).

8.3 Why the Solow Model Fails as a Business Cycle Model

A natural first attempt is to use the Solow model we already know as a theory of fluctuations. Suppose TFP A_t fluctuates around a constant level A^*—say, $A_t \in \{A^* e^{\varepsilon}, A^* e^{-\varepsilon}\}$ for some small $\varepsilon > 0$. In the Solow diagram, this shifts the saving curve $sA_t f(k_t)$ up and down, causing the steady-state capital stock to oscillate between two values. The economy would indeed exhibit cyclical behavior in k_t and y_t.

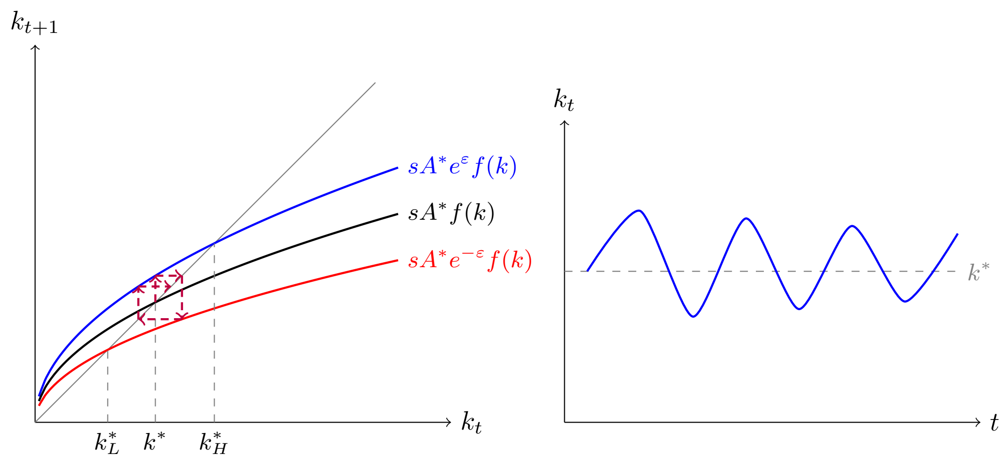

However, the Solow model imposes a **constant saving rate**: $I_t = sY_t$. This immediately implies

$$\ln I_t = \ln s + \ln Y_t.$$

Taking variances of both sides:

$$\text{Var}[\ln I_t] = \text{Var}[\ln Y_t] + \underbrace{\text{Var}[\ln s]}_{=0},$$

so $\text{Var}[\ln I_t] = \text{Var}[\ln Y_t]$. That is, investment is *exactly as cyclical* as output.

Remark.

This directly contradicts **Fact 1**. In the data, the variance of cyclical investment is several times larger than that of cyclical output. The source of the failure is clear: the Solow model treats the saving rate as an exogenous constant, so agents cannot choose to save more in booms and less in recessions. To generate excess volatility of investment, we need a model in which the saving (equivalently, consumption-investment) decision is made optimally by forward-looking agents. This is the motivation for the **Real Business Cycle (RBC)** model, which replaces the mechanical saving rule with a fully microfounded dynamic optimization problem.

8.4 An RBC Model with Fixed Labor

Having established that the Solow model fails to match the key business cycle facts due to its constant saving rate, we now build a model in which the consumption-saving decision is made optimally by a forward-looking planner. This is the **Real Business Cycle (RBC)** model, originating from Kydland and Prescott (1982). In this simplified version, we fix labor supply and focus on how stochastic productivity shocks, combined with optimal intertemporal choice, generate cyclical fluctuations that are qualitatively consistent with the data.

8.4.1 Environment

- **Time:** Discrete, $t = 0, 1, 2, \ldots$
- **Preferences:** A representative planner maximizes expected discounted lifetime utility:

$$U\left(\{c_t\}_{t=0}^{\infty}\right) = \mathbb{E}\left(\sum_{t=0}^{\infty} \beta^t u(c_t)\right),$$

  where $\beta \in (0,1)$ is the discount factor and $u(\cdot)$ is a strictly concave, increasing period utility function.

- **Production:** The aggregate production function exhibits constant returns to scale: $Y_t = A_t F(K_t, L)$. Since labor L is fixed (normalized to 1), we can write everything in per-capita terms. Define $k_t \equiv K_t/L$ and $f(k_t) \equiv F(K_t/L,\, L/L) = F(k_t, 1)$. Then output per capita is:

$$y_t = e^{a_t} f(k_t),$$

  where $a_t \equiv \ln A_t$ is log-TFP. The stochastic process $\{a_t\}$ is the sole source of randomness in the model.

- **Stochastic Process for Technology:** The log-TFP process $\{a_t\}_{t=0}^{\infty}$ follows a **Markov process** with a time-invariant conditional distribution $\pi(a_{t+1} \mid a_t)$.
- **Resource Constraint:** In every period,

$$c_t + k_{t+1} = e^{a_t} f(k_t) + (1-\delta) k_t, \quad \forall\, t \geq 0,$$

  where $\delta \in (0,1)$ is the depreciation rate.

Remark (Timing and Information).

The timing within each period is as follows. At the beginning of period t, the planner enters with capital k_t (chosen last period) and observes the realization of a_t, drawn from $\pi(a_t \mid a_{t-1})$. After observing a_t, the planner chooses consumption c_t and next-period capital k_{t+1}. Then at $t+1$, the new shock a_{t+1} is realized, and the process repeats.

Crucially, the planner's decision at time t can depend on the *current* state (a_t, k_t) but not on future realizations of a. This is why we look for a policy function of the form $k_{t+1} = g(a_t, k_t)$.

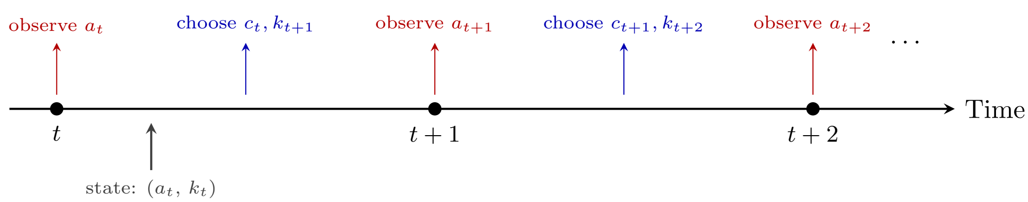

8.4.2 The Planner's Problem

By the principle of optimality, the planner's problem can be written recursively. The **value function** $V(a, k)$ satisfies the Bellman equation:

$$V(a,k) = \max_{\substack{c>0 \\ k'\geq 0}} \{u(c) + \beta\,\mathbb{E}(V(a', k') \mid a)\}$$
$$\text{s.t.} \quad c + k' = e^a f(k) + (1-\delta)k.$$

The associated **policy function** is:

$$g(a,k) = \arg\max_{\substack{c>0 \\ k'\geq 0}} \{u(c) + \beta\,\mathbb{E}(V(a', k') \mid a)\} \quad \text{s.t.} \quad c + k' = e^a f(k) + (1-\delta)k.$$

Unlike the deterministic Solow model, the saving rate here is *endogenous*: the planner optimally chooses how much to consume and how much to invest, conditional on the current state of technology. When productivity is temporarily high, the planner may choose to save a larger fraction of output (investing for the future), and when productivity is low, the planner may dissave. This endogenous response is precisely what allows the model to generate investment volatility in excess of output volatility.

Remark (Why a Planner? The Welfare Theorems).

We solve a representative-planner problem rather than a decentralized market with households, firms, and prices. This is a deliberate shortcut justified by the **First Welfare Theorem**: in this environment—complete markets, no externalities, no taxes, perfect competition—any competitive equilibrium is Pareto efficient, and conversely, any Pareto-efficient allocation can be supported as a competitive equilibrium with appropriate prices (Second Welfare Theorem). Solving the planner's problem therefore yields the same allocation $\{c_t, k_{t+1}\}$ as a fully decentralized RBC economy.

The competitive prices "hidden inside" the planner's solution can be recovered ex post:

$$\underbrace{w_t = e^{a_t} f'(k_t) \cdot (\text{stuff})}_{\text{real wage = MPL}}, \qquad \underbrace{r_t = e^{a_t} f'(k_t) - \delta}_{\text{real interest rate = net MPK}}, \qquad \underbrace{\beta \frac{u'(c_{t+1})}{u'(c_t)}}_{\text{stochastic discount factor}}.$$

> This trick—solve the easier planner problem, recover prices later—is pervasive in modern macro. It *fails* once frictions are introduced (incomplete markets, distortionary taxes, sticky prices), at which point one must work with the decentralized equilibrium directly. New Keynesian models break the equivalence on purpose—that is what makes monetary policy non-trivial there.

8.4.3 Numerical Solution by Value Function Iteration

To solve the model concretely, consider a tractable special case in which technology takes only two values:

$$a_t \in \{a_H, a_L\}, \quad a_H > a_L,$$

with transitions governed by a 2×2 Markov transition matrix:

$$\Pi = \begin{pmatrix} \pi_{HH} & \pi_{HL} \\ \pi_{LH} & \pi_{LL} \end{pmatrix},$$

where $\pi_{ij} = \Pr(a_{t+1} = a_j \mid a_t = a_i)$ and each row sums to 1.

Remark.

> Persistence in the productivity process (i.e., π_{HH} and π_{LL} being large) is important for generating realistic business cycles. If shocks were purely i.i.d., the economy would revert to its mean too quickly, producing cycles that are too short and too shallow relative to the data.

The Bellman equation generally does not admit a closed-form solution, so we solve it numerically. The key idea is to **discretize** the capital state space and iterate on the value function until convergence.

To begin with, we discretize the capital space into a grid of N points:

$$\mathcal{K} = \{k^1, k^2, \ldots, k^N\}.$$

The objects we seek are:

- **Value function:** An $N \times 2$ matrix $\mathbf{V} = [\mathbf{V}_H, \mathbf{V}_L]$, where $V_H^n \equiv V(a_H, k^n)$ and $V_L^n \equiv V(a_L, k^n)$.
- **Policy function:** An $N \times 2$ matrix $\mathbf{g} = [\mathbf{g}_H, \mathbf{g}_L]$, where $g_H^n \equiv g(a_H, k^n)$ and $g_L^n \equiv g(a_L, k^n)$.

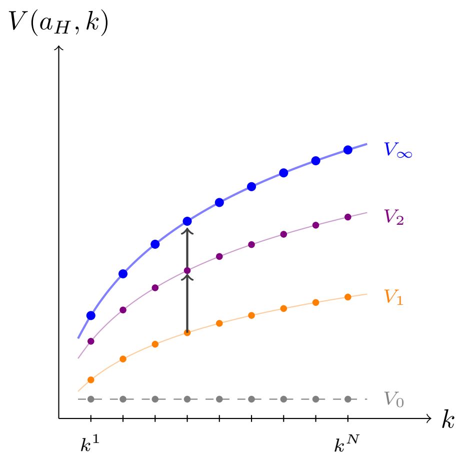

Value Function Iteration

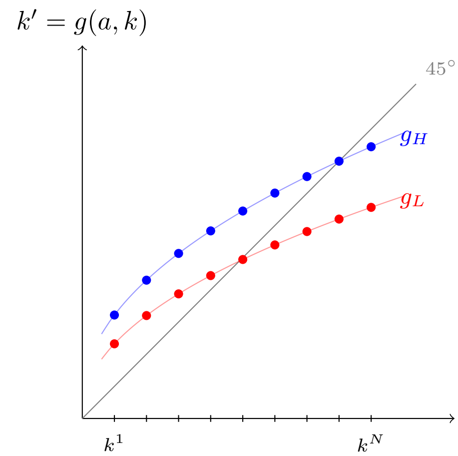

Policy Functions

Value Function Iteration Algorithm for the RBC Model

Step 1: Initialize. Guess an initial value function $\mathbf{V}_0 = [\mathbf{V}_{H,0}, \mathbf{V}_{L,0}]$. A natural starting point is $\mathbf{V}_0 = \mathbf{0}$ (the $N \times 2$ zero matrix). Set the iteration counter $\texttt{nI} = 0$.

Step 2: Update. For each grid point k^n and each technology state, compute the updated value by solving a finite maximization problem. For the high state:

$$V^n_{H,\texttt{nI}+1} = \max_{n' \in \{1,\dots,N\}} \left\{ u\left[e^{a_H} f(k^n) + (1-\delta)k^n - k^{n'}\right] + \beta\left(\pi_{HH}\, V^{n'}_{H,\texttt{nI}} + \pi_{HL}\, V^{n'}_{L,\texttt{nI}}\right)\right\},$$

and analogously for the low state:

$$V^n_{L,\texttt{nI}+1} = \max_{n' \in \{1,\dots,N\}} \left\{ u\left[e^{a_L} f(k^n) + (1-\delta)k^n - k^{n'}\right] + \beta\left(\pi_{LH}\, V^{n'}_{H,\texttt{nI}} + \pi_{LL}\, V^{n'}_{L,\texttt{nI}}\right)\right\}.$$

Record the optimizer $n'_H(n)$ and $n'_L(n)$ at each grid point. Set $\texttt{nI} \leftarrow \texttt{nI}+1$.

Step 3: Check convergence. If $\|\mathbf{V}_{\texttt{nI}} - \mathbf{V}_{\texttt{nI}-1}\| < \varepsilon$ for a small tolerance $\varepsilon > 0$, stop. Otherwise, return to Step 2.

Step 4: Extract the policy function. Using the optimizers from the final iteration, construct:

$$g^n_H = k^{n^*_H(n)}, \qquad g^n_L = k^{n^*_L(n)}, \qquad \text{for } n = 1, \dots, N.$$

Remark (Properties of the Solution).

The converged solution has two intuitive properties:

- **Policy functions:** $g_H(k) > g_L(k)$ for all k—when productivity is high, the planner invests more (chooses higher k'). Both policy functions lie below the 45-degree line for large k (capital eventually decays if no new investment is made) and above it for small k.

- **Value functions:** $V_H(k) > V_L(k)$ for all k—the planner is strictly better off when productivity is high, holding capital fixed. Both are increasing and concave in k.

8.4.4 Simulation and Evaluating the Model

With the policy function in hand, we can simulate the model economy and check whether it reproduces the key business cycle facts.

Step 1: Draw shocks. Simulate a sequence $\{a_t\}_{t=0}^{T}$ (e.g., $T = 100$) from the Markov chain Π.

Step 2: Generate capital path. Starting from an initial k_0, iterate forward using the policy function:

$$k_{t+1} = g(a_t, k_t), \quad t = 0, 1, \ldots, T.$$

Step 3: Recover all variables. For each t:

$$y_t = e^{a_t} f(k_t), \qquad I_t = k_{t+1} - (1-\delta)k_t, \qquad c_t = y_t - I_t.$$

Step 4: Compute business cycle statistics. To compare the cyclicality of variables that differ in scale (output is much larger than investment in levels), we use the **coefficient of variation**—the standard deviation divided by the mean:

$$\mathrm{CV}(x_t) \equiv \frac{\sqrt{\mathrm{Var}[x_t]}}{\bar{x}} = \sqrt{\mathrm{Var}\left[\frac{x_t}{\bar{x}}\right]}.$$

This is a unit-free measure of volatility: it tells us the *percentage* fluctuation of a variable around its mean. For instance, $\mathrm{CV}(I_t) = 0.10$ means investment fluctuates by roughly $\pm 10\%$ around its average level.

The RBC model predicts:

$$\mathrm{CV}(I_t) > \mathrm{CV}(y_t) > \mathrm{CV}(c_t),$$

i.e., **investment is more volatile than output, and consumption is smoother than output**.

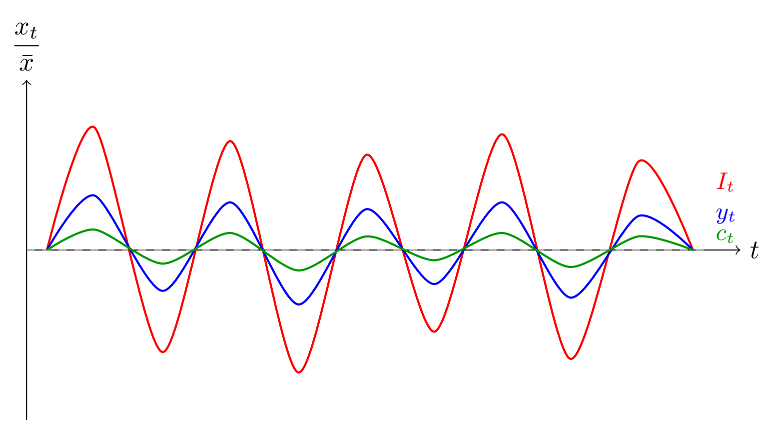

Impulse Response: Anatomy of a Single Shock

A complementary way to look at the model's mechanics is to ask: starting from the deterministic steady state, what happens if a single positive productivity shock hits at $t = 0$ and then dies out at the AR(1) rate ρ? The path of $\{a_t, y_t, c_t, I_t, k_t\}$ in response to this one-time shock is called the **impulse response function (IRF)**.

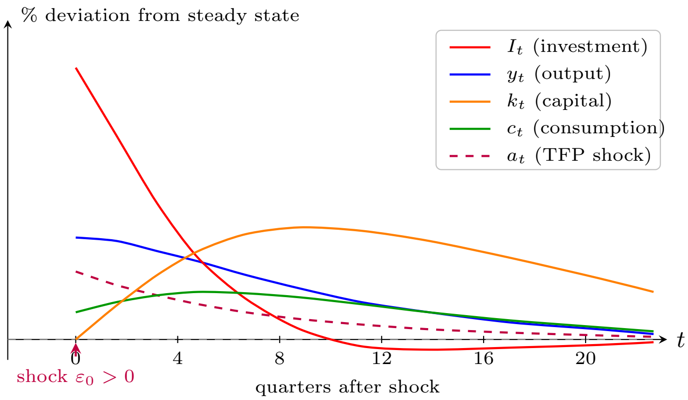

Three features of this picture are worth flagging:

- **Investment overshoots on impact.** I_t jumps far above y_t on impact—the planner front-loads investment to take advantage of temporarily high productivity.
- **Consumption is smooth and hump-shaped.** c_t rises modestly and *keeps rising* for several quarters even though productivity is already decaying. This is consumption smoothing: the agent spreads the windfall across many periods.
- **Capital propagates the shock.** Even after a_t has nearly returned to zero, the capital stock k_t remains elevated for many quarters, keeping output above trend. This is the *internal propagation mechanism* of the RBC model: even a one-period iid shock would generate persistent fluctuations through capital accumulation. Persistence in a_t amplifies this further.

Remark (Internal vs. External Propagation).

> A common critique of basic RBC models (Cogley and Nason, 1995) is that their internal propagation is actually quite weak: most of the persistence in simulated y_t comes from the assumed persistence of a_t itself, not from endogenous capital dynamics. If you set $\rho = 0$ (iid shocks), the model produces virtually no autocorrelation in output. Subsequent literature added richer propagation channels—variable capacity utilization, labor adjustment costs, habit formation—to address this.

Remark (Why the RBC Model Succeeds Where Solow Failed).

> The key difference is the endogenous saving rate. When a positive productivity shock hits ($a_t = a_H$), output rises. The forward-looking planner recognizes that productivity is *persistent but mean-reverting*, so current income is temporarily high relative to

permanent income. By the logic of consumption smoothing, the planner saves a disproportionately large share of the windfall, causing investment to spike. Conversely, when $a_t = a_L$, the planner dissaves to smooth consumption, and investment drops sharply.[a] The result is that investment inherits the full force of the productivity shock *plus* an amplification from the intertemporal substitution motive, while consumption is deliberately smoothed. This is precisely what generates $\text{Var}[I_t] > \text{Var}[y_t] > \text{Var}[c_t]$—matching Fact 1.

[a]Here "dissaving" means that the saving rate falls below its long-run average; net investment $I_t = k_{t+1} - (1-\delta)k_t$ remains positive in typical calibrations, but it is much smaller than in booms.

8.4.5 General Model: Beyond Two States

The two-state example in the previous subsection captures the essential logic of the RBC model, but it is too coarse to produce quantitatively realistic business cycles. We now generalize to allow log-TFP to follow a **continuous AR(1) process**:

$$a' = \rho\, a + \varepsilon, \qquad \varepsilon \overset{\text{iid}}{\sim} N(0, \sigma_\varepsilon^2),$$

where $\rho \in (0,1)$ controls the persistence of productivity shocks and σ_ε governs their volatility.

Remark (Standard Calibration).

To make the model quantitative, parameters are typically calibrated to match long-run U.S. data at a quarterly frequency:

Parameter	Meaning	Value	Target / source
β	discount factor	0.99	annual real interest rate $\approx 4\%$
δ	depreciation rate	0.025	$\sim 10\%$ per year
α	capital share in $F(K,L) = K^\alpha L^{1-\alpha}$	1/3	U.S. capital income share
σ	CRRA / inverse IES	1–2	micro consumption studies
ρ	persistence of log-TFP	0.95	autocorrelation of Solow residual
σ_ε	std. dev. of TFP innovation	0.007	matches $\text{Var}\left[\ln y_t^C\right]$

The values for ρ and σ_ε come directly from estimating an AR(1) on the (HP-filtered) Solow residual; the rest are pinned down by long-run averages or independent micro evidence. With these numbers, the simulated fixed-labor economy reproduces $\text{CV}(I)/\text{CV}(y) \approx 3$ and $\text{CV}(c)/\text{CV}(y) \approx 0.5$—roughly matching the U.S. postwar data. (The variable-labor extension introduced later in this chapter adds one more parameter, the Frisch elasticity ν, calibrated to $\nu \approx 4$ to match the cyclicality of total hours.)

The Bellman equation is unchanged:

$$V(a,k) = \max_{\substack{c>0 \\ k' \geq 0}} \{u(c) + \beta\, \mathbb{E}(V(a',k') \mid a)\}$$
$$\text{s.t.} \quad c + k' = e^a f(k) + (1-\delta)k,$$

but the conditional expectation now integrates over a continuous normal distribution rather than summing over two states.

The solution procedure proceeds in five steps.

Step 1: Approximate the AR(1) Process

Since we cannot work with a continuous state space for a on a computer, we approximate the AR(1) process by a finite-state Markov chain. Choose a grid of M points:

$$\vec{a} = [a^1, a^2, \dots, a^M],$$

together with an $M \times M$ transition matrix Π, where $\pi_{ij} = \Pr\big(a_{t+1} = a^j \mid a_t = a^i\big)$.

Remark (The Tauchen Method).

The standard method for constructing this discrete approximation is due to Tauchen (1986). The idea is to choose the grid points $\{a^1, \dots, a^M\}$ to span the ergodic distribution of the AR(1) process (typically covering ± 3 unconditional standard deviations), and then to set π_{ij} equal to the probability that the AR(1) process lands in the bin around a^j given that it starts at a^i—computed using the conditional normal CDF $\Phi\Big(\frac{a^j + \Delta/2 - \rho\, a^i}{\sigma_\varepsilon}\Big) - \Phi\Big(\frac{a^j - \Delta/2 - \rho\, a^i}{\sigma_\varepsilon}\Big)$, where Δ is the grid spacing. As $M \to \infty$, the Markov chain approximation converges to the true AR(1) process. In practice, M between 5 and 15 is usually sufficient.

Geometrically, the joint density of (a_t, a_{t+1}) under the AR(1) process is a bivariate normal whose contours are ellipses tilted along the line $a_{t+1} = \rho\, a_t$. Tauchen overlays an $M \times M$ rectangular grid on this density and assigns π_{ij} to be the conditional mass in column j given row i:

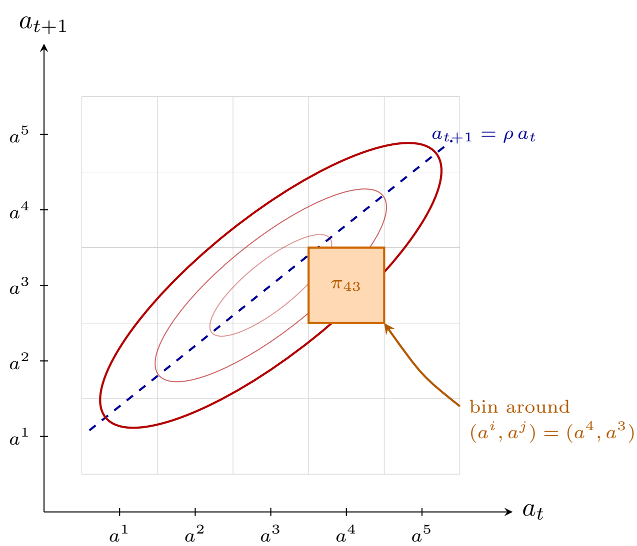

The red ellipses are level sets of the joint density $f(a_t, a_{t+1})$; the dashed line is the conditional mean $\mathbb{E}(a_{t+1} \mid a_t) = \rho a_t$. The Tauchen probability π_{ij} is the conditional mass of the orange bin centered at (a^i, a^j) given $a_t = a^i$—i.e., the integral of the conditional density along the row $a_t = a^i$ over the bin's vertical extent.

Step 2: Discretize the Capital Space

Choose a grid of N points for capital:

$$\mathcal{K} = [k^1, k^2, \ldots, k^N].$$

Step 3: Represent Value and Policy Functions as Matrices

The value function $V(a,k)$ is now approximated by an $N \times M$ matrix:

$$\mathbf{V} = [\mathbf{V}_1, \mathbf{V}_2, \ldots, \mathbf{V}_M] = \begin{pmatrix} V_1^1 & V_2^1 & \cdots & V_M^1 \\ V_1^2 & V_2^2 & \cdots & V_M^2 \\ \vdots & & \ddots & \vdots \\ V_1^N & V_2^N & \cdots & V_M^N \end{pmatrix},$$

where $V_m^n \equiv V(a^m, k^n)$. Similarly, the policy function is an $N \times M$ matrix:

$$\mathbf{g} = [\mathbf{g}_1, \mathbf{g}_2, \ldots, \mathbf{g}_M] = \begin{pmatrix} g_1^1 & g_2^1 & \cdots & g_M^1 \\ g_1^2 & g_2^2 & \cdots & g_M^2 \\ \vdots & & \ddots & \vdots \\ g_1^N & g_2^N & \cdots & g_M^N \end{pmatrix},$$

where $g_m^n \equiv g(a^m, k^n)$ gives the optimal next-period capital when the current state is (a^m, k^n).

Step 4: Solve by Value Function Iteration

Apply the same iterative algorithm as in the two-state case, now generalized to M technology states. For each grid point (a^m, k^n), the update rule is:

$$V^n_{m,\mathtt{nI}+1} = \max_{n' \in \{1,\ldots,N\}} \left\{ u\Big[e^{a^m} f(k^n) + (1-\delta)k^n - k^{n'}\Big] + \beta \sum_{j=1}^{M} \pi_{mj}\, V^{n'}_{j,\mathtt{nI}} \right\}.$$

Iterate until $\|\mathbf{V}_{\mathtt{nI}+1} - \mathbf{V}_{\mathtt{nI}}\| < \varepsilon$, then extract the policy function from the final-iteration optimizers.

The converged policy functions have the same qualitative properties as in the two-state case, but now there are M curves $g_1, g_2, \ldots, g_M$ in the (k, k') plane, one for each technology state, with $g_m(k) < g_{m'}(k)$ whenever $a^m < a^{m'}$.

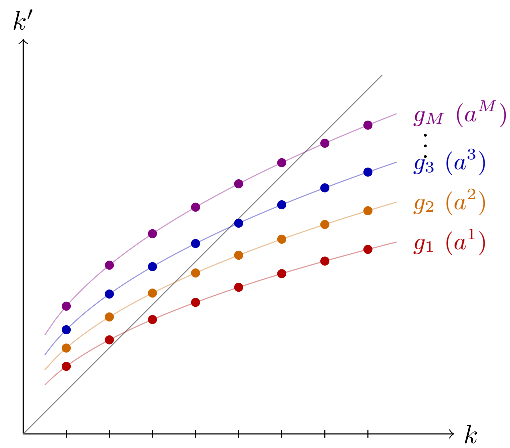

Step 5: Simulate

Given initial conditions k_0 and a_0, simulate $\{a_t\}_{t=0}^T$ from the Markov chain Π, then iterate:

$$k_{t+1} = g(a_t, k_t), \qquad y_t = e^{a_t} f(k_t), \qquad I_t = k_{t+1} - (1-\delta)k_t, \qquad c_t = y_t - I_t.$$

The simulated series $\{k_{t+1}, c_t, y_t, I_t\}$ can then be used to compute business cycle statistics and compare the model's predictions against the data.

8.5 An RBC Model with Endogenous Labor

The fixed-labor RBC model can match Facts 1 and 2 (investment volatility and pro-cyclical productivity), but it is silent on Fact 3 (cyclicality of total hours). To address all three facts simultaneously, we now extend the model to allow the planner to choose **labor supply** L in addition to consumption and investment.

8.5.1 The Planner's Problem

Since labor L is now a choice variable, we write the problem in aggregate terms (C, K, L) rather than per-capita terms.[1] The Bellman equation in aggregate terms is:

$$\begin{aligned} V(a, K) = \max_{C,\,L,\,K'} & \{u(C, L) + \beta\,\mathbb{E}(V(a', K') \mid a)\} \\ \text{s.t.} \quad & C + K' = e^a F(K, L) + (1-\delta)K, \\ & C > 0, \quad K' \geq 0, \quad L \in (0, 1). \end{aligned}$$

The period utility $u(C, L)$ now depends on both consumption and labor. We assume:

$$u_C > 0, \quad u_{CC} < 0, \quad u_L < 0, \quad u_{LL} < 0.$$

[1]There are two reasons. First, a technical one: when L varies over time, per-capita variables like $k_t = K_t/L_t$ have the endogenous choice L_t in the denominator, making the transformation messy (though not impossible). Second, and more fundamentally: the whole point of introducing endogenous labor is to model the *fluctuation of aggregate hours* L and test whether it matches Fact 3. If we divided through by L and worked in per-capita terms, the variable L would be normalized away from the Bellman equation, and we would lose the ability to study its cyclical behavior. We need L to appear explicitly as a choice variable.

The first pair says the agent likes consumption with diminishing marginal utility. The second pair says the agent dislikes working, and the marginal disutility of labor is *increasing*: each additional hour of work is more painful than the last.

8.5.2 Characterizing the Solution: Two Euler Equations

We derive the optimality conditions in three steps:

Step 1: Take FOCs of the Bellman equation with respect to C, L, and K'.

Step 2: Apply the envelope theorem to express the derivative of the value function in terms of period-t marginal utilities.

Step 3: Substitute the envelope condition into the FOCs to eliminate the value function. This yields two Euler equations.

Inter-temporal Euler Equation (EE_K)

The FOC with respect to K' and the envelope condition together give:

$$u_C(C,L) = \beta\,\mathbb{E}\left(u_C(C',L')\Big(\underbrace{e^{a'}F_K(K',L')}_{\text{MPK}'} + 1 - \delta\Big) \,\middle|\, a\right) \tag{EE_K}$$

This is the standard consumption Euler equation: the marginal cost of saving one more unit today (left side: foregone marginal utility of consumption) equals the expected marginal benefit tomorrow (right side: the additional unit of capital earns the marginal product of capital $e^{a'}F_K(K',L')$ net of depreciation, valued at tomorrow's marginal utility).

Intra-temporal Euler Equation (EE_L)

The FOC with respect to L gives:

$$-u_L(C,L) = u_C(C,L)\cdot \underbrace{e^a F_L(K,L)}_{\text{MPL}} \tag{EE_L}$$

This is a **static** condition—it involves only time-t variables. It says the marginal disutility of working one more hour (left side) must equal the marginal utility of the extra consumption that hour produces (right side: the marginal product of labor, converted to utility units by multiplying by u_C).[2]

Remark.

> Note that (EE_K) is an *inter-temporal* condition (linking t and $t+1$), while (EE_L) is an *intra-temporal* condition (a within-period trade-off between leisure and consumption). Together with the resource constraint, they fully characterize the optimal allocation.

[2] MPL $\equiv e^a F_L(K,L)$ is the *marginal product of labor*: the additional output produced by one more unit of labor, holding capital fixed. MPK $\equiv e^a F_K(K,L)$ is the *marginal product of capital*: the additional output from one more unit of capital, holding labor fixed. In a competitive equilibrium, MPL equals the real wage and MPK equals the rental rate of capital.

8.5.3 Parameterization and the Frisch Elasticity

To make the model quantitative, we adopt a standard parameterization that separates consumption utility from labor disutility:

$$u(C,L) = \frac{C^{1-\sigma}}{1-\sigma} - B\frac{L^{1+\frac{1}{\nu}}}{1+\frac{1}{\nu}},$$

where $\sigma > 0$ is the coefficient of relative risk aversion (governing the curvature of consumption utility), $B > 0$ scales the disutility of labor, and $\nu > 0$ is a parameter that controls the labor supply elasticity.

Under this specification, the partial derivatives are:

$$u_C = C^{-\sigma}, \qquad u_L = -BL^{1/\nu}.$$

Plugging into (EE_L):

$$BL^{1/\nu} = C^{-\sigma} \cdot \text{MPL},$$

which, after rearranging, gives:

$$L = \left(\frac{1}{B}\right)^{\nu} C^{-\sigma\nu} \cdot \text{MPL}^{\nu}.$$

Taking logs:

$$\ln L = \text{const.} - \sigma\nu \ln C + \nu \ln \text{MPL}.$$

Definition 8.4: Frisch Elasticity of Labor Supply

The **Frisch elasticity** measures the responsiveness of labor supply to changes in the wage (marginal product of labor), *holding the marginal utility of wealth constant*:

$$\text{Frisch Elasticity} \equiv \left.\frac{\partial \ln L}{\partial \ln \text{MPL}}\right|_{u_C=\text{const.}} = \nu.$$

A higher ν means that hours worked respond more strongly to changes in the wage. This elasticity is the key parameter governing how much labor fluctuates over the business cycle.

Remark.

- The condition $u_C = \text{const.}$ isolates the *intertemporal substitution* channel. To see why, consider a worker deciding how many hours to work today vs. tomorrow. Suppose wages are temporarily high today but will return to normal tomorrow. The worker faces two competing effects:

  Substitution effect: "Wages are high today, so each hour of work buys more consumption. I should work more today and take leisure tomorrow." This effect increases current labor supply.

  Wealth effect: "I'm richer overall because of the high wage. Since I'm richer, I can afford more leisure." This effect decreases current labor supply.

The Frisch elasticity captures *only* the substitution effect by holding u_C constant. Here is why this shuts down the wealth channel. When wages rise permanently, the agent is richer in a lifetime sense. A richer agent consumes more (C rises), which means $u_C = C^{-\sigma}$ *falls*. Looking back at the labor supply equation $\ln L = \text{const.} - \sigma\nu \ln C + \nu \ln \text{MPL}$, we see that higher C (lower u_C) reduces L—this is precisely the wealth effect ("I'm richer, so I consume more leisure"). By holding u_C constant, we freeze C at its current level and ask: if only the wage changes but the agent's wealth position does not, how much does labor supply respond? This isolates the pure substitution motive. If the wage increase is truly temporary (as business cycle shocks are), the wealth effect is small (one good quarter barely changes lifetime income), so the Frisch elasticity is the relevant measure. If the wage increase were permanent, the wealth effect would be large, and the total (Marshallian) labor supply elasticity—which combines both effects—would be smaller than the Frisch elasticity.

- For any variable x, $\mathrm{d}\ln x = \mathrm{d}x/x$ is approximately the percentage change. So the Frisch elasticity $\nu = \partial \ln L / \partial \ln \text{MPL}$ is literally "the percent change in hours worked per one percent change in the wage"—exactly the standard notion of elasticity from intermediate micro, applied to the specific context of labor supply. What makes it "Frisch" is the ceteris paribus condition: we hold the marginal utility of wealth fixed, which isolates the pure substitution effect of a wage change.

8.5.4 Discussion: The Frisch Elasticity Puzzle

The model with endogenous labor faces a serious quantitative tension when we try to match all three business cycle facts simultaneously.

What the Model Needs

Consider the labor market response to a positive productivity shock (a increases). The shock raises the MPL, which shifts the labor demand curve outward from L_L^D to L_H^D. How much equilibrium hours change depends on the slope of the labor supply curve, which is governed by the Frisch elasticity ν.

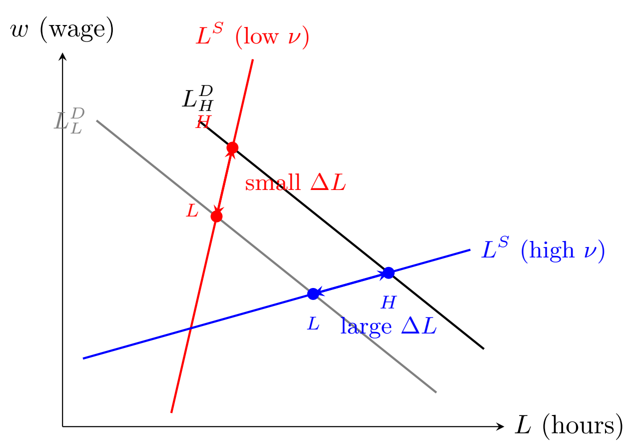

As the diagram shows, a steep labor supply curve (low ν, red) implies that the same demand shift produces only a small change in equilibrium hours—most of the adjustment shows up in wages. A flat supply curve (high ν, blue) instead generates a large change in hours with only a modest wage response.

To match **Fact 3** (total hours are as cyclical as output), the model needs the latter scenario: labor supply must be highly elastic. Quantitatively, matching both Fact 2 and Fact 3 simultaneously requires a Frisch elasticity of:

$$\hat{\nu}_{\text{RBC}} > 4.$$

What Micro Data Say

The micro labor literature estimates the Frisch elasticity using individual-level data—measuring how much a given worker adjusts hours in response to temporary wage changes (e.g., overtime premia, tax holidays). The consensus from this literature is:

$$\hat{\nu}_{\text{micro}} \ll 1, \quad \text{typically } 0.1\text{–}0.2.$$

That is, individual workers barely adjust their hours when wages change temporarily. This is strikingly at odds with the $\hat{\nu}_{\text{RBC}} > 4$ required by the RBC model—a discrepancy of more than an order of magnitude.

Potential Resolutions

There are two main avenues for reconciling the high macro elasticity with the low micro elasticity:

1. Aggregation via the Extensive Margin. Recall from Fact 3 that roughly 2/3 of the cyclical variation in total hours comes from the extensive margin (workers entering and leaving employment), not the intensive margin (hours per worker). Micro estimates of the Frisch elasticity capture only the *intensive margin*—how much an already-employed worker adjusts hours. But at the macro level, what matters is the *aggregate* response of total hours $\sum_i h_i$, which includes the extensive margin response of workers who move in and out of employment entirely. Even if each individual has a low intensive-margin elasticity, the aggregate elasticity can be much larger if the extensive margin is responsive.

Bottom line: *extensive-margin aggregation closes roughly* 1/3 *of the gap between* $\hat{\nu}_{micro}$ *and* $\hat{\nu}_{RBC}$*—partially, but not all.*

Remark (Why the Extensive Margin Is Missed by Micro Estimates).

> The logic is as follows. Micro studies estimate the Frisch elasticity by regressing an individual's hours on their wage, using workers who are *continuously employed*. This captures the intensive margin: how an employed person adjusts hours at the margin. But during a recession, many workers do not reduce hours—they lose their jobs entirely (or during a boom, previously non-employed workers enter the labor force). These discrete 0-to-1 or 1-to-0 transitions constitute the extensive margin. Since a worker who exits employment has no "hours vs. wage" data point, the extensive margin response is invisible in micro panel regressions.
>
> At the aggregate level, total hours = (number of workers) × (hours per worker). The micro Frisch elasticity governs only the second term. Empirically, the first term is highly

> responsive to aggregate conditions: during the Great Recession, average hours per worker fell by only about 1%, while the number of employed workers fell by roughly 6%. This is not just a theoretical possibility—it is a well-documented pattern (see Fact 3: 2/3 of hours variation is extensive margin). Workers who remain employed show little change in their individual hours (consistent with $\hat{\nu}_{\text{micro}} \approx 0.1$–$0.2$), but the aggregate labor input swings dramatically because many workers transition between employment and non-employment. This is why the *macro* Frisch elasticity (governing the response of total hours) can far exceed the *micro* Frisch elasticity. However, accounting for this channel closes roughly 1/3 of the gap, leaving a substantial residual discrepancy.

Remark (Hansen (1985) and Indivisible Labor).

> A classic theoretical device for closing the gap is Hansen's (1985) **indivisible labor** model. Suppose individual workers face a fixed cost of going to work (commuting, child-care, training), so each worker chooses only between $L_i = 0$ and $L_i = \bar{h}$ (e.g., 40 hours per week). With this $0/\bar{h}$ corner-solution structure, the standard Frisch elasticity at the individual level is undefined or zero—the worker doesn't smoothly slide along a labor supply curve.
>
> But suppose households can write *employment lotteries*: each household commits in advance to send a fraction $\phi \in [0, 1]$ of its members to work, and the lottery pools the resulting consumption risk. Then the relevant decision variable is ϕ, which is continuous, and the aggregate hours-supply schedule becomes *linear* in the wage. The implied aggregate Frisch elasticity is *infinite*, even though every individual has zero intensive-margin elasticity.
>
> The key insight: **indivisibilities at the micro level can produce a high macro elasticity through the extensive margin alone**, with no need to assume that any individual is highly elastic. This is essentially a sharper, theoretically tractable version of the aggregation argument above. The catch is the lottery assumption itself—real households cannot literally insure against employment risk this way (markets are incomplete), so Hansen's setup is best read as a parable rather than a literal description.

2. Non-competitive or Frictional Labor Markets. The RBC model assumes a perfectly competitive, frictionless labor market in which workers are always on their labor supply curve. In particular, the wage always equals the MPL, and employment is determined by the intersection of supply and demand. This is what forces the model to need a high Frisch elasticity: the *only* way to generate large hours fluctuations in a competitive market is to have a flat supply curve.

But real labor markets are not frictionless. Two important departures:

- **Wage stickiness.** Suppose wages cannot adjust freely in the short run—they are "sticky" due to long-term contracts, social norms against wage cuts, or institutional constraints.

  Consider first a *negative* productivity shock. The MPL falls below the prevailing wage $\bar{w}$. The firm now faces a choice: cut wages for all workers to match the new MPL, or lay off some workers while keeping the remaining workers' wages unchanged. In practice, firms overwhelmingly choose layoffs over across-the-board wage cuts.[3] Workers who keep their jobs face the same wage as before, so their individual labor supply barely changes—consistent with a low micro Frisch elasticity. The entire drop in aggregate hours comes from workers who are laid off (extensive margin), a discrete jump that has nothing to do with anyone sliding along a labor supply curve.

The logic is symmetric in *booms*. When a positive shock raises the MPL above $\bar{w}$, firms find it profitable to hire additional workers at the existing wage rather than raise wages for current employees. Each new hire generates surplus since MPL $> \bar{w}$. The result is a large increase in employment with little movement in wages or incumbent workers' hours.

In both directions, wage rigidity channels aggregate fluctuations into the extensive margin, generating large swings in total hours without requiring a high Frisch elasticity.

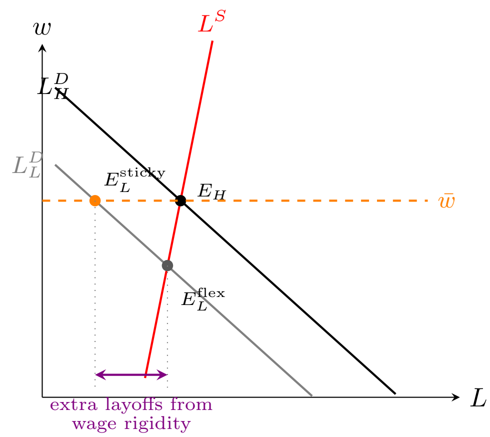

Under flexible wages, a negative shock moves the equilibrium from E_H to E_L^{flex}: both wages and hours fall modestly. Under sticky wages, the wage stays at $\bar{w}$, so the firm slides along its demand curve to E_L^{sticky}, laying off all workers whose MPL falls below the rigid wage. The violet arrow marks the additional employment loss caused purely by wage rigidity.

Remark (A Caveat on Identification).

> An important limitation of the wage stickiness explanation is that it is difficult to test directly. If wages truly do not move, then we never observe the counterfactual flexible-wage equilibrium, and we cannot identify the slope of the labor supply curve from aggregate data. The observed variation in hours is driven entirely by demand shifts along a rigid wage, so the data are equally consistent with *any* value of the micro Frisch elasticity. In this sense, sticky wages can "explain" any pattern of hours fluctuations, making the theory hard to falsify. However, there is independent micro-level evidence that wages are indeed sticky (e.g., Bewley, 1999; Barattieri et al., 2014), which lends external validity to this channel.

[3] This is a robust empirical finding. Bewley (1999) documents through extensive interviews that managers strongly resist nominal wage cuts, citing concerns about worker morale, effort, and retention.

- **Monopsony power.** In a perfectly competitive labor market, firms are price-takers: they hire until MPL $= w$. Under monopsony, a single firm (or a firm with market power) faces an upward-sloping labor supply curve—to hire more workers, it must raise the wage for *all* workers. This creates a wedge: the *marginal cost of labor* (MCL) exceeds the wage, because hiring one more worker raises the wage bill for all existing workers.

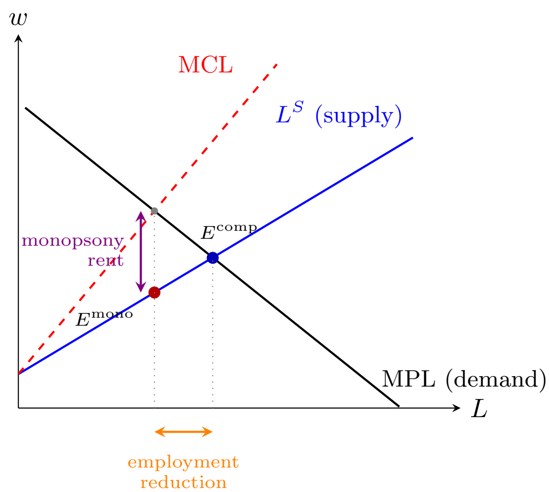

In a competitive market, firms hire until MPL $= w$, reaching E^{comp}. A monopsonist recognizes that hiring an extra worker raises the wage for *all* workers, so the marginal cost of labor (MCL, red dashed) lies above the supply curve. The monopsonist hires where MCL $=$ MPL, reaching E^{mono} with fewer workers (orange gap) paid a lower wage. The violet arrow shows the **monopsony rent**: the wedge between what the last worker produces and what the firm pays. When productivity shifts, the monopsonist adjusts employment along the MCL $=$ MPL margin—governed by the firm's market power, not by workers' labor supply elasticity. This breaks the tight link between the Frisch elasticity and aggregate hours fluctuations.

These insights motivate **New Keynesian** models, which embed nominal rigidities (sticky prices, sticky wages) into the RBC framework. In such models, aggregate hours can fluctuate substantially over the business cycle—driven by demand-side forces and wage rigidity—without requiring an implausibly high Frisch elasticity. The elasticity ν remains a property of individual preferences (a "result" of the utility function parameterization), but in a New Keynesian world, the *equilibrium response of hours* is no longer pinned down by ν alone—it also depends on the degree of wage stickiness and market power.

8.6 Where the RBC Model Falls Short

The RBC model is a remarkable success: a single technology shock, propagated through optimal intertemporal choice, qualitatively reproduces all three key business cycle facts and provides a tractable laboratory for quantitative policy analysis. But it has well-known weaknesses, and recognizing them is what motivated the next generation of models.

- **What is a "technology shock"?** The model attributes recessions to negative shocks to a_t. Taken literally, this requires that society periodically *forgets* how to produce things—an interpretation many economists find implausible. In practice, the Solow residual conflates true technology, capacity utilization, labor effort, markups, and measurement error. Basu et al. (2006) show that purified technology shocks behave very differently from raw Solow residuals.
- **The Galí (1999) puzzle.** Using long-run identification in a structural VAR, Galí finds that positive technology shocks *reduce* hours worked on impact—the opposite of the RBC prediction. This finding has been contested but remains an influential challenge to the basic model.
- **The labor wedge.** Even within the model's own framework, the intratemporal Euler equation $-u_L/u_C = \text{MPL}$ fails badly in the data: there is a large, time-varying gap (the "labor wedge") between the marginal rate of substitution and the marginal product. This wedge accounts for a large share of cyclical fluctuations in hours (Chari et al., 2007) and points squarely at frictions the RBC model abstracts from.
- **No role for monetary policy.** The model contains no money, no nominal prices, and no interest rate that the central bank could set. Yet decades of empirical evidence suggest monetary policy has real effects. RBC simply has nothing to say about this.
- **Welfare cost of cycles.** If cycles are the optimal response of a frictionless economy, then stabilization policy is at best useless and at worst harmful—and the welfare cost of fluctuations is tiny (Lucas, 1987, calculated $\sim 0.05\%$ of consumption). Most policymakers and many economists find this conclusion hard to swallow, especially after large recessions.

The standard pedagogical arc is to use the RBC model as a baseline—to learn how to write down a stochastic dynamic general-equilibrium model, solve it numerically, and confront it with the data—and then to introduce frictions one at a time: nominal rigidities (New Keynesian), incomplete markets and household heterogeneity (HANK), search and matching in the labor market (DMP), financial frictions (Bernanke–Gertler), and so on. Each extension preserves the RBC scaffolding but breaks one of its strong assumptions to address one of the shortcomings above.

Remark (Chapter Summary).

- **Three key business-cycle facts.** Investment is more volatile than output; productivity (Solow residual) is pro-cyclical; total hours are about as cyclical as output, with the extensive margin contributing roughly 2/3.
- **Why Solow fails.** The exogenous saving rate ties $\text{Var}[(]\ln I) = \text{Var}[(]\ln Y)$ exactly, contradicting Fact 1. The RBC model fixes this by replacing the saving rule with optimal intertemporal choice.
- **Numerical solution.** Discretize the AR(1) shock via Tauchen and the capital state via a grid; iterate the Bellman equation. Persistence in TFP combined with capital propagation generates serial correlation in output.
- **The endogenous-labor extension.** Adds an intratemporal Euler equation linking labor disutility to MPL. Matching Fact 3 requires Frisch elasticity $\nu > 4$, far above

the micro consensus of 0.1–0.2. Indivisible labor (Hansen 1985) and extensive-margin aggregation partially close the gap; sticky wages and search frictions complete it.

- **Robust shortcomings.** Technology-shock interpretation is contestable; Galí (1999)'s sign reversal in the labor response is hard to reconcile; the labor wedge is large; the model has no role for monetary policy; the welfare cost of cycles is implausibly small.
- **What comes next.** Chapter 9 turns the RBC model into a measurement device (Business Cycle Accounting); Chapter 10 refines its consumption side; modern New Keynesian and HANK models add the frictions the RBC model abstracts from.

Chapter 9

Business Cycle Accounting

Remark (Notation in This Chapter).

Symbol	Meaning
a_t	efficiency wedge (log-TFP shock, retains the role from Ch. 8)
τ_{Lt}	labor wedge: distortion of the consumption-leisure margin
τ_{It}	investment wedge: distortion of the saving margin
g_t	Government-spending wedge
$\vec{z}_t$	Stacked wedge state vector $(a_t, \tau_{It}, \tau_{Lt}, g_t)$
P, Σ	AR(1) transition matrix and innovation covariance for $\vec{z}_t$
$T(\vec{z}, K)$	Lump-sum rebate of tax revenue
EE_K, EE_L	Distorted inter- and intra-temporal Euler equations

The previous chapter built the canonical RBC model and showed that a single technology shock can qualitatively reproduce the three key business cycle facts. But the qualitative match is far from a quantitative success: at standard parameter values the model under-predicts the cyclicality of hours, mis-times the comovement of consumption and investment, and—most pointedly—requires an implausibly high Frisch elasticity to fit aggregate hours. Real economies clearly violate the frictionless competitive benchmark. The natural next question is: *by how much, and along which margin?*

Business Cycle Accounting (BCA), developed by Chari et al. (2007, henceforth **CKM**), provides a systematic answer. The strategy is to take the RBC model not as a literal description of the economy but as a *measurement device*. We deliberately insert "wedges" into the equilibrium conditions of an otherwise standard RBC model, treat each wedge as an unobserved stochastic process, and let the data tell us which wedges fluctuated when. The resulting decomposition does not say *which* friction caused a recession, but it tells us *which margin* of the frictionless model the recession violated most. That, in turn, sharply narrows the set of structural theories worth taking seriously.

9.1 The Concept: Wedges as Tax-Equivalent Distortions

Suppose we observe in the data that the marginal product of capital is higher than the user cost predicted by the consumption Euler equation. The competitive RBC model predicts these two should be equal. The gap between them is, by definition, a **wedge**: a violation of the optimality condition that we can interpret *as if* an economic agent had been levying a tax on the relevant transaction.

Remark (Wedges are Tax-Equivalents, Not Literal Taxes).

> The wedges in BCA are not interpreted as actual government taxes. They are reduced-form summaries of *any* friction—sticky prices, market power, financial constraints, search costs, information frictions—that drives a wedge between the marginal rate of substitution and the marginal rate of transformation. CKM prove a series of **equivalence theorems** showing that detailed models with very different frictions (e.g., a sticky-price model and a financial-friction model) generate observationally identical paths for $\{C, I, L, Y\}$ to a prototype RBC model with appropriate wedges. BCA exploits this equivalence to back out which margins are distorted, leaving the question of *which* structural friction is responsible for downstream theoretical work.

The Harberger Triangle

The geometry is the standard one from public finance. Consider a single intertemporal margin: the household equates the marginal cost of saving (foregone consumption today, $u'(C)$) to the discounted marginal benefit (extra consumption tomorrow, $\beta u'(C')\,\text{MPK}'$). In the data, however, we observe a gap: at the realized K^{data}, the marginal benefit of one more unit of capital exceeds its marginal cost. The gap is the wedge:

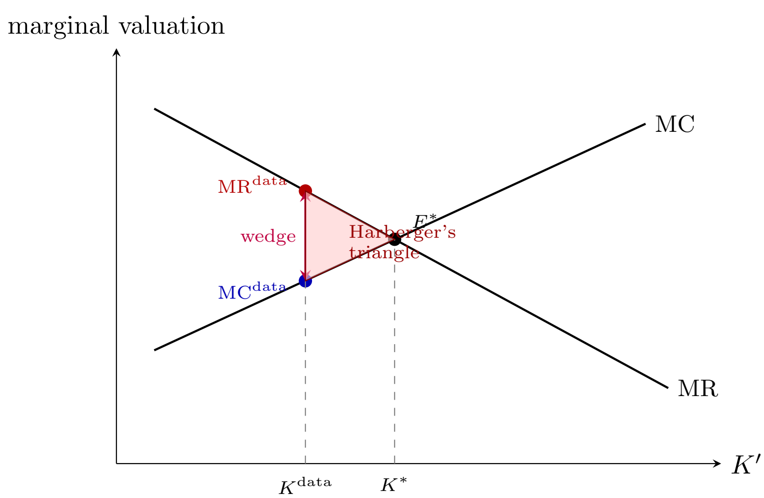

The vertical gap at K^{data} is the marginal wedge—the size of the distortion expressed in the units of marginal valuation. Translated into a tax rate, even a wedge of a few percentage points can represent a substantial tax-equivalent. The shaded triangle is the standard **Harberger triangle**: it measures the deadweight loss generated by the wedge, equal to the area between the MR and MC curves over the distorted region $[K^{\text{data}}, K^*]$.

Remark (Reading the Wedge off an Euler Equation).

In the frictionless RBC model the consumption Euler equation is

$$u'(C_t) = \beta\,\mathbb{E}(u'(C_{t+1})\,\text{MPK}_{t+1})\,.$$

Plug in the data $\{C_t, K_t\}$. Generically, the LHS will not equal the RHS—there will be a residual. The BCA strategy is to attribute the entire residual to a single multiplicative "tax" τ_I on the cost of investment:

$$(1+\tau_{I,t})\,u'(C_t) = \beta\,\mathbb{E}(u'(C_{t+1})\,[\,F'_K + (1+\tau_{I,t+1})(1-\delta)\,])\,.$$

The same logic applies to the intratemporal labor-leisure margin and yields a labor wedge τ_L. The wedges are not estimated structural parameters; they are *whatever* residuals make the equilibrium conditions hold given the observed data. They are by construction reduced-form—which is precisely why they are so informative as a diagnostic device.

9.2 Setup: A Prototype Economy with Stochastic Wedges

The starting point of BCA is a **decentralized RBC model** augmented with three reduced-form distortions:

- **Efficiency wedge** a_t: the standard log-TFP shock from the RBC model. Captures shocks to total factor productivity, but also any unmodeled friction that distorts the production function (e.g., variable capacity utilization, sectoral misallocation).
- **Labor wedge** τ_{Lt}: a tax on labor income that distorts the household's intratemporal labor-leisure choice. Captures any friction that drives a gap between the marginal rate of substitution between consumption and leisure and the marginal product of labor (sticky wages, monopsony, search frictions, distortionary income taxes).
- **Investment wedge** τ_{It}: a tax on the cost of investment that distorts the household's intertemporal saving choice. Captures any friction that drives a gap between the marginal cost and marginal benefit of capital accumulation (financial constraints, collateral requirements, intermediation spreads, capital income taxes).

A fourth process,

- **Government consumption** g_t: stochastic government purchases of goods and services,

is included not so much because government spending is the central object of interest, but because we need a fourth shock for the model to deliver a determinate equilibrium for the four observables $\{C_t, I_t, L_t, Y_t\}$. Without it, the system would be over-identified.

Remark (Why We Need Four Shocks).

> BCA aims to invert the model: given four observed series $\{C, I, L, Y\}$, recover four exogenous processes $\{a, \tau_L, \tau_I, g\}$. For this inversion to be well-posed, the number of shocks must equal the number of observables. Three would leave one observable over-determined; five would leave one shock under-identified. The government-spending wedge g_t is the natural fourth process: it is a real, observable category in the national accounts, but in BCA it functions primarily as the residual that lets the linear system close.

Each of the four processes is assumed to follow an AR(1):

$$z_{t+1}^{(j)} = \rho_j \, z_t^{(j)} + \varepsilon_{t+1}^{(j)}, \qquad \varepsilon^{(j)} \overset{\text{iid}}{\sim} N(0, \sigma_j^2),$$

for $j \in \{a, \tau_L, \tau_I, g\}$. Stacking, the state vector is

$$\boxed{\vec{z}_t \equiv (a_t,\ \tau_{It},\ \tau_{Lt},\ g_t)}$$

with vector AR(1) representation $\vec{z}_{t+1} = P\vec{z}_t + \vec{\varepsilon}_{t+1}$, where P is a 4×4 matrix and $\vec{\varepsilon}_{t+1} \sim N(\vec{0}, \Sigma)$. The matrix P and covariance Σ are the parameters BCA *estimates*; the path of $\vec{z}_t$ is the object BCA *recovers*.

9.3 Recursive Competitive Equilibrium

Because BCA inserts wedges into a *decentralized* environment (rather than working with a planner's problem as in the previous chapter), we must specify the household problem, the firm problem, market clearing, and the government budget separately.

9.3.1 Household

The representative household, taking prices and policy as given, chooses consumption, labor, and capital to maximize lifetime utility. In recursive form:

$$V(\vec{z}, K) = \max_{C,\, L,\, K'} \{u(C, L) + \beta\, \mathbb{E}(V(\vec{z}', K') \mid \vec{z})\}$$

subject to the after-tax budget constraint

$$\boxed{C + (1 + \tau_I)\,[\,K' - (1 - \delta)K\,] \;=\; (1 - \tau_L)\, w(\vec{z}, K)\, L \;+\; r(\vec{z}, K)\, K \;+\; T(\vec{z}, K)}$$

where T is a lump-sum rebate of tax revenue (which we will pin down through the government budget constraint).

The interpretation of each tax:

- $(1 - \tau_L)$ multiplies labor income wL: the household receives only a fraction $(1 - \tau_L)$ of every wage dollar earned.
- $(1 + \tau_I)$ multiplies net investment $I = K' - (1 - \delta)K$: every unit of new capital costs the household $(1 + \tau_I)$ units of consumption goods.

9.3.2 Firm

A competitive representative firm, taking the wage and rental rate as given, hires labor and capital each period:

$$\max_{K,\,L}\; e^{a}\,F(K,L)\;-\;w(\vec{z},K)\,L\;-\;r(\vec{z},K)\,K.$$

The firm operates statically (no capital accumulation decision; that is in the household's problem). The first-order conditions give the standard factor-price equations

$$w(\vec{z},K)=e^{a}F_L(K,L),\qquad r(\vec{z},K)=e^{a}F_K(K,L).$$

9.3.3 Market Clearing and Government Budget

Goods market clearing requires

$$C\;+\;\underbrace{[\,K'-(1-\delta)K\,]}_{=I}\;+G\;=\;e^{a}F(K,L).$$

The government budget balances tax revenue against spending plus transfers:

$$G+T=\tau_L\,w(\vec{z},K)\,L\;+\;\tau_I\,[\,K'-(1-\delta)K\,].$$

Remark (Why Lump-Sum Rebates?).

> The lump-sum transfer T is a modeling convenience: it ensures that the wedges τ_L,τ_I act as *pure* distortionary taxes (no income effect from financing government spending). Without it, an increase in τ_L would simultaneously distort labor supply *and* make the household poorer, mixing two channels. With the rebate, raising τ_L only changes relative prices on the margin—a cleaner experiment, and the one CKM use throughout.

9.3.4 Definition

Definition 9.1: Recursive Competitive Equilibrium

A **recursive competitive equilibrium (RCE)** for the prototype economy consists of

- a value function $V(\vec{z}, K)$ and policy functions $g(\vec{z}, K) = (C(\vec{z}, K),\ L(\vec{z}, K),\ K'(\vec{z}, K),\ Y(\vec{z}, K))$,
- pricing functions $w(\vec{z}, K)$ and $r(\vec{z}, K)$,

such that

1. Given prices and policy, the household's problem is solved by V and g;
2. Given prices, the firm's problem is solved at the chosen (K, L);
3. The goods market clears;
4. The government budget balances.

9.4 Wedges in the Equilibrium Conditions

The equilibrium of the prototype economy is conveniently summarized by two distorted Euler equations and the resource constraint. These three equations, together with the AR(1) processes for $\vec{z}_t$, fully characterize the model's predictions for $\{C_t, I_t, L_t, Y_t\}$.

Intratemporal Euler Equation (the labor margin)

The household's first-order condition for labor, combined with the firm's FOC $w = e^a F_L$, yields:

$$\boxed{-u_L(C, L) \;=\; (1 - \tau_L)\, u_C(C, L) \cdot \underbrace{e^a F_L(K, L)}_{\text{MPL}}} \tag{EE_L}$$

Comparing to the frictionless RBC condition $-u_L = u_C \cdot \text{MPL}$, the wedge $(1-\tau_L)$ enters multiplicatively on the MPL side. Whenever the data show households working "too little" relative to what the wage would justify in a competitive market, BCA registers this as a positive labor wedge $\tau_L > 0$.

Intertemporal Euler Equation (the saving margin)

The household's first-order condition for K' is:

$$\boxed{(1 + \tau_{I,t})\, u_C(C_t, L_t) \;=\; \beta\, \mathbb{E}(u_C(C_{t+1}, L_{t+1})\, [\, e^{a_{t+1}} F_K(K_{t+1}, L_{t+1}) \;+\; (1 + \tau_{I,t+1})(1 - \delta)\,] \,|\, \vec{z}_t)} \tag{EE_K}$$

Two features deserve attention. First, τ_I enters multiplicatively on the cost side ($1+\tau_{I,t}$ on the LHS) and on the resale value of undepreciated capital tomorrow ($1 + \tau_{I,t+1}$ inside the expectation). Second, the labor wedge τ_L does *not* appear in (EE_K)—only τ_I does.

This separation is what makes the two wedges identifiable from one another: each distorts a different margin.

Remark (Which Margin, Which Wedge?).

> A useful mnemonic:
>
> - The **labor wedge** τ_L distorts the *within-period* trade-off between consumption and leisure. It shows up in (EE_L).
> - The **investment wedge** τ_I distorts the *across-period* trade-off between consumption today and consumption tomorrow. It shows up in (EE_K).
> - The **efficiency wedge** a_t distorts the production function itself. It shows up wherever F does, which is in both Euler equations and the resource constraint.
> - The **government wedge** g_t enters only through the resource constraint (a demand component that competes with C and I).

9.5 Estimation: How CKM Recover the Wedges

The model has now been transformed into a state-space system: four observables $\{C_t, I_t, L_t, Y_t\}$, four latent stochastic processes $\vec{z}_t = (a_t, \tau_{It}, \tau_{Lt}, g_t)$, and equilibrium conditions linking them. The estimation strategy is canonical for such systems.

The Chari et al. (2007) Estimation Procedure

Step 1: Calibrate non-stochastic parameters. Fix $\beta, \delta, \alpha, \sigma, \nu$, and the form of the utility and production functions to standard values from the RBC literature.

Step 2: Take HP-filtered data. Detrend $\{C_t, I_t, L_t, Y_t\}$ and work with their cyclical components.

Step 3: Log-linearize the equilibrium conditions around the deterministic steady state. The resulting system is linear in $\{\hat{C}_t, \hat{I}_t, \hat{L}_t, \hat{Y}_t\}$ and $\{\hat{a}_t, \hat{\tau}_{It}, \hat{\tau}_{Lt}, \hat{g}_t\}$.

Step 4: Estimate P and Σ by maximum likelihood. Treat the AR(1) parameters (P, Σ) as unknowns. Use the Kalman filter to compute the likelihood of the observed data given the model and parameters. Maximize over (P, Σ).

Step 5: Recover the wedge series. With $(\hat{P}, \hat{\Sigma})$ in hand, run the Kalman smoother to back out the most likely path of $\vec{z}_t$ given the entire observed sample. This yields the estimated wedges $\{\hat{a}_t, \hat{\tau}_{Lt}, \hat{\tau}_{It}, \hat{g}_t\}_t$.

Remark (Why Maximum Likelihood and Not OLS?).

Naively, one might try to compute $\hat{\tau}_{Lt}$ at each date by plugging $\{C_t, L_t, K_t\}$ into the intratemporal Euler equation and solving for τ_{Lt}. This works in principle but ignores the cross-equation restrictions: in a rational-expectations equilibrium, the household's choice of C_t depends on its forecast of *all four* future wedges, and that forecast depends on the joint AR(1) parameters (P, Σ). The MLE/Kalman approach uses the full system to extract wedges efficiently and to estimate (P, Σ) jointly. It also gives standard errors and lets us run counterfactuals (e.g., "what would output have done if only the efficiency wedge had moved?").

9.6 What BCA Has Found

The headline result of Chari et al. (2007), and of the large BCA literature that followed, is striking and remarkably robust across countries and time periods:

The labor wedge does most of the work.

More specifically:

- During major U.S. recessions—the Great Depression of 1929–39, the 1982 recession, the Great Recession of 2008–09—the bulk of the cyclical decline in output is attributable to movements in the **efficiency wedge** a_t and the **labor wedge** τ_{Lt}, with the labor wedge typically explaining the larger share of hours and a comparable share of output.
- The **investment wedge** τ_{It} is surprisingly small. Even in episodes that look superficially like "credit crunches" (the Great Depression, 2008), the investment wedge plays a minor role. This was one of the most controversial findings of CKM, since it suggested that financial-friction models—then ascendant in macro—were missing the main story.
- The **government-spending wedge** g_t explains very little of cyclical fluctuations in $\{C, I, L, Y\}$, consistent with the prior view that fiscal shocks are not the dominant driver of postwar U.S. business cycles.

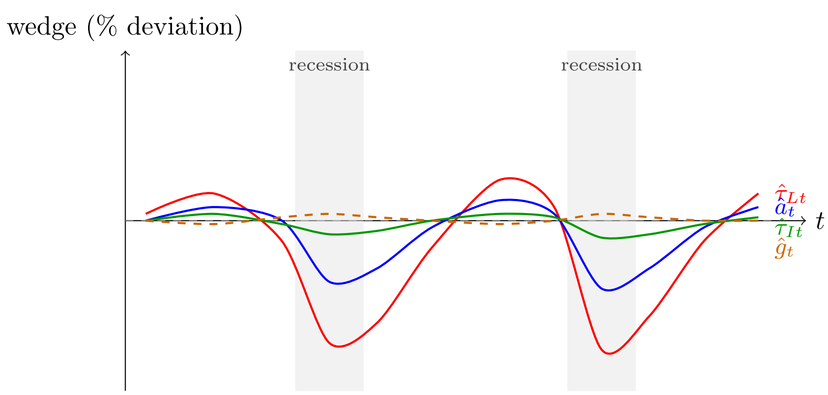

The figure above is a stylized rendering of the CKM result: in each NBER-dated recession (gray bars), the labor wedge falls sharply, the efficiency wedge falls moderately, the investment wedge barely moves, and the government wedge is essentially flat.

Remark (The Ranking Matters More Than the Levels).

> The exact percentages CKM report depend on the calibration, the sample, the detrending method, and the specification of the production function. But across many sensitivity checks—and across many countries (Kehoe and Prescott 2007 collected studies of dozens of recessions worldwide)—the qualitative ranking $\hat{\tau}_L \succ \hat{a} \succ \hat{\tau}_I \succ \hat{g}$ in importance is remarkably stable. This is the durable lesson of BCA, even if any single point estimate can be debated.

Implication: Where Theory Should Focus

If the labor wedge is what fluctuates most over the cycle, then the structural friction we are looking for is one that (i) primarily distorts the consumption-leisure margin, and (ii) varies systematically over the cycle. Candidate stories that pass this filter include:

- **Sticky wages.** Nominal wage rigidity, as in New Keynesian models, generates a time-varying gap between MRS and MPL. When aggregate demand falls, real wages remain too high, so employers hire too little and the labor wedge widens.
- **Search and matching frictions.** In Diamond–Mortensen–Pissarides labor markets, wages reflect a Nash bargain between workers and firms rather than the marginal product, and the bargain shifts with labor-market tightness over the cycle.
- **Distortionary taxes and transfers.** Cyclical changes in marginal tax rates, unemployment insurance, or means-tested transfers can mechanically generate a cyclical labor wedge.
- **Monopsony / monopoly markups.** Time-varying market power directly produces a wedge between MPL and the wage paid.

Conversely, models whose primary mechanism distorts the *investment* margin (e.g., pure financial friction models in the spirit of Bernanke–Gertler) face a quantitative challenge: BCA says the investment wedge does not move much, so a friction operating exclusively through that margin will not explain the bulk of cyclical fluctuations. (This does not rule out financial frictions altogether: a financial friction might amplify shocks via labor-market channels, in which case BCA would attribute the action to τ_L, not τ_I.)

9.7 Limitations and Critiques

BCA is a measurement device, not a structural model. Several caveats are important to bear in mind.

- **Wedges are not structural.** BCA tells us *which margin* is distorted, not *what causes* the distortion. Different deep theories can produce identical wedge patterns (this is the equivalence theorem). Moving from a BCA finding to a structural theory still requires additional identifying assumptions.
- **Specification dependence.** Christiano and Davis (2006) showed that BCA results can be sensitive to the assumed period utility function and to whether some wedges are

restricted to enter as multiplicative taxes vs. as additive shifters. The qualitative ranking is robust; the precise decomposition is not.

- **Linearization.** Standard BCA log-linearizes around the steady state. In severe recessions or near a binding constraint, the linearization may misattribute nonlinear dynamics to one wedge or another. Recent work has extended BCA with nonlinear filters.
- **The labor wedge is not a clean object.** A growing literature (Karabarbounis 2014; Bils et al. 2018) has shown that decomposing τ_L further—into a household side (the gap between MRS and the wage) and a firm side (the gap between the wage and MPL)—reveals that most of the cyclical labor wedge is on the household side. This has reignited the debate about whether the labor wedge reflects sticky wages, mismeasurement of the consumption-leisure trade-off, or something else entirely.

Remark (BCA in Perspective).

> The deeper contribution of BCA is methodological. It taught macroeconomists to ask, before writing down a detailed structural model: *which margin of the frictionless benchmark is the one I am trying to explain?* A model whose mechanism operates on a margin that the data say is undistorted is, from a quantitative perspective, a non-starter—no matter how elegant. Conversely, any model that successfully matches a margin BCA flags as cyclically distorted is at least in the right neighborhood. This has shaped the agenda of macro for two decades: from sticky-price models that endogenize the labor wedge, to search models that microfound it, to recent work on heterogeneous-agent New Keynesian (HANK) models that connect the labor wedge to inequality and the marginal propensity to consume.

Remark (Chapter Summary).

> - **BCA repurposes the RBC model.** Insert four reduced-form "wedges" (efficiency, labor, investment, government) into the equilibrium conditions and back them out from the data via Kalman filtering of a linearized state-space system.
> - **Wedges are tax-equivalents, not literal taxes.** Equivalence theorems (CKM 2007) show that wedges with appropriate stochastic processes can replicate paths of $\{C, I, L, Y\}$ generated by detailed structural models with very different micro frictions.
> - **The labor wedge does most of the work.** Across decades, countries, and detrending choices, $\hat{\tau}_L$ accounts for the bulk of cyclical variation in major recessions; the investment wedge $\hat{\tau}_I$ is small. The government wedge contributes little.
> - **Implications for theory.** Models whose primary mechanism distorts the consumption-leisure margin (sticky wages, search and matching, monopsony) survive the BCA filter. Pure financial-friction models that operate exclusively through τ_I do not.
> - **Limitations.** Wedges are reduced-form, not structural; results depend on log-linearization; the labor wedge can be further decomposed into household-side and firm-side components, and most of the cyclicality is on the household side.

Chapter 10

Consumption and Saving

Remark (Notation in This Chapter).

Symbol	Meaning
Y_t	Period-t income (deterministic in PIH section, stochastic in RWH section)
A_t	Period-t asset holdings
PI	Permanent income $\equiv A_0 + \sum_{t\geq 0} Y_t/(1+r)^t$
$\phi(r)$	Annuity factor, $\phi(r) = (1 - 1/(1+r))/(1 - (1/(1+r))^{T+1})$
Y_i^P, Y_i^T	Permanent and transitory components of cross-sectional income
$\hat{b}$	Cross-sectional regression coefficient (Keynes), $= \text{Var}(Y^P)/\text{Var}(Y)$
σ	Coefficient of relative risk aversion under CRRA preferences
$X_t = \beta^t u'(C_t)$	Discounted marginal utility (martingale under RWH)
ε_{t+1}	Consumption innovation $C_{t+1} - C_t$ in the Hall test
λ	Hand-to-mouth fraction in Campbell–Mankiw
PFD	Alaska Permanent Fund Dividend
$\hat{\alpha}$	Excess-sensitivity coefficient (Hsieh, Kueng)
s^t	History of income realizations through t (RWH stochastic setup)

The previous chapters spent considerable effort *computing* aggregate consumption and saving in stochastic, general-equilibrium models. We now step back and ask the underlying behavioral question: **how does a single household decide how much to consume in each period given a stream of income?** The answer that emerges—once we take optimization and forward-looking behavior seriously—is sharper and more counterintuitive than the simple rule "consume what you earn." This is the core insight of the **Permanent Income Hypothesis (PIH)** of Modigliani and Brumberg (1954) and Friedman (1957), together with its stochastic refinement, Hall's (1978) **Random Walk Hypothesis (RWH)**.

We develop the theory in two stages. The first half of the chapter treats the household problem under *certainty*, deriving the PIH in its cleanest form. The second half reintroduces uncertainty and derives Hall's random-walk implication. The final section previews the modern empirical literature that tests both.

10.1 The Big Picture: Consumption & Saving, Data vs. Theory

The long arc of macro consumption theory pits two views against one another:

- **Keynes (1936):** consumption is mainly a function of *current* income. The marginal propensity to consume (MPC) out of an extra dollar today is positive and stable, around 0.6–0.8.
- **Friedman (1957) PIH:** consumption is a function of *permanent income*—the annuity value of total lifetime resources. Transitory income shocks barely affect consumption; only permanent shifts matter.

These two views make sharply different predictions about how households respond to a tax cut, a year-end bonus, a winning lottery ticket, or a temporary recession. Most of the modern consumption literature can be read as adjudicating between them, modifying both, or measuring how far reality lies from each.

Remark (Outline of the Topic).

> Following the standard pedagogical order:
>
> - **Without uncertainty:** the *Permanent Income Hypothesis* (next section).
> - **With uncertainty:** Hall's *Random Walk Hypothesis*, which adds Euler-equation testable implications (the section after that).
> - **Modern empirical literature:** natural-experiment-based tests of MPC, Excess Sensitivity, excess smoothness, precautionary saving (final section).

10.2 The Permanent Income Hypothesis: No Uncertainty

We begin with the deterministic case. Removing uncertainty is not because it is realistic—it is not—but because it is the cleanest setting in which to derive the central PIH theorem. Uncertainty will be layered on top in the next section.

10.2.1 Setup

The key idea, in one sentence:

utility maximization + a perfect (state-contingent) financial market ⇒ consumption depends only on the net present value of lifetime wealth.

Removing uncertainty from the problem makes the "state-contingent" qualifier vacuous—there is only one state—so what remains is just frictionless intertemporal borrowing and lending at a fixed interest rate.

Simplifying Assumptions

- **Finite horizon:** $t = 0, 1, 2, \ldots, T$.
- **Representative household** with a known, exogenous income process $\{Y_t\}_{t=0}^{T}$ (no labor-supply choice; income is just an endowment stream).
- **Borrowing and lending** at a constant interest rate r, with full commitment (the household always pays back).
- **Initial assets** A_0 given.

The Household's Problem

The household chooses $\{C_t, A_{t+1}\}_{t=0}^{T}$ to maximize

$$\boxed{\max_{\{C_t,\,A_{t+1}\}_{t=0}^{T}} \sum_{t=0}^{T} \beta^t\, u(C_t)}$$

subject to the period budget constraint

$$\frac{A_{t+1}}{1+r} + C_t \;=\; Y_t + A_t, \qquad \forall\, t = 0, 1, \ldots, T,$$

non-negativity $C_t \geq 0$, A_0 given, and the terminal condition

$$A_{T+1} \geq 0 \qquad \text{(natural borrowing limit)}.$$

Remark (Reading the Period Budget).

> The form $\frac{A_{t+1}}{1+r} + C_t = Y_t + A_t$ uses the convention that A_t is the household's wealth at the *start* of period t, after last period's interest has accrued. The household then earns income Y_t, consumes C_t, and the remainder is saved at end-of-period in an account that returns $(1+r)$ next period. Equivalently: the saving made today is worth A_{t+1} tomorrow, so its date-t value is $A_{t+1}/(1+r)$. The terminal condition $A_{T+1} \geq 0$ rules out dying in debt; together with the household's strict preference for consumption ($u' > 0$), it will bind with equality at the optimum—the household leaves no money on the table.

Remark ("Source of the Trade-off").

> Why is there a saving decision at all? Because resources moved across time pay a price/reward: a dollar saved today returns $(1+r)$ tomorrow, and a dollar borrowed today must be paid back as $(1+r)$ tomorrow. Marginal utility tomorrow is discounted by $\beta < 1$. The household balances these two forces—the **intertemporal trade-off**—which is exactly what the Euler equation will pin down.

10.2.2 From Period Budgets to the Intertemporal Budget Constraint

The $T+1$ period constraints can be collapsed into a single **intertemporal budget constraint** (IBC) by repeated substitution. Rearrange the period-t constraint:

$$\frac{A_{t+1}}{1+r} - A_t \;=\; Y_t - C_t.$$

Divide by $(1+r)^t$:

$$\frac{A_{t+1}}{(1+r)^{t+1}} - \frac{A_t}{(1+r)^t} \;=\; \frac{Y_t - C_t}{(1+r)^t}.$$

Sum from $t=0$ to $t=T$. The LHS telescopes:

$$\frac{A_{T+1}}{(1+r)^{T+1}} - A_0 \;=\; \sum_{t=0}^{T} \frac{Y_t - C_t}{(1+r)^t}.$$

Rearrange and impose $A_{T+1} \geq 0$:

$$\boxed{\underbrace{\sum_{t=0}^{T} \frac{C_t}{(1+r)^t}}_{\text{PV of lifetime consumption}} \quad\leq\quad \underbrace{A_0 \;+\; \sum_{t=0}^{T} \frac{Y_t}{(1+r)^t}}_{\text{PV of lifetime wealth (“Permanent Income” PI)}}\;.} \qquad \text{(IBC)}$$

At the optimum (since utility is strictly increasing in C_t), the IBC binds with equality.

Definition 10.1: Permanent Income

The **Permanent Income** of a household is the present value of its total lifetime resources:

$$\text{PI} \;\equiv\; A_0 \;+\; \sum_{t=0}^{T} \frac{Y_t}{(1+r)^t}.$$

It is a single scalar that summarizes everything the household knows about its budget—initial wealth plus the discounted sum of all current and future labor income.

Remark (The IBC Collapses Time).

The move from $T+1$ period constraints to a single IBC is more than a notational convenience. It says: *from a budgeting standpoint, the household does not face a sequence of separate problems, one per period—it faces a single lifetime allocation problem.* The household can shift consumption arbitrarily across dates as long as the present-value sum is respected. This is what the perfect-financial-market assumption buys you: complete flexibility to move resources through time.

10.2.3 Characterizing the Solution: The PIH Theorem

Maximizing a strictly concave objective subject to a linear constraint gives a clean characterization.

Proposition (Necessary and Sufficient Conditions)

A consumption sequence $\{C_t\}_{t=0}^T$ solves the household's problem if and only if it satisfies:

(1) The **intertemporal budget constraint** with equality:

$$\sum_{t=0}^{T} \frac{C_t}{(1+r)^t} \;=\; \text{PI}.$$

(2) The **Euler equation** at every $t = 0, 1, \ldots, T-1$:

$$u'(C_t) \;=\; \beta\,(1+r)\,u'(C_{t+1}). \tag{EE}$$

The Euler equation is the standard intertemporal optimality condition: at the optimum, the marginal utility of a dollar consumed today equals the discounted marginal utility of that same dollar saved and consumed tomorrow.

The PIH Theorem

The IBC + EE system has a striking implication.

The Permanent Income Hypothesis

The optimal consumption path $\{C_t\}_{t=0}^T$ depends on the income stream $\{Y_t\}_{t=0}^T$ **only through Permanent Income**.

Why? Because the entire income stream enters the system only via the IBC, and only via the scalar PI. The Euler equation does not involve Y_t at all. So two different income streams that share the same PI lead to the *identical* consumption path—even though their year-by-year profiles can be wildly different.

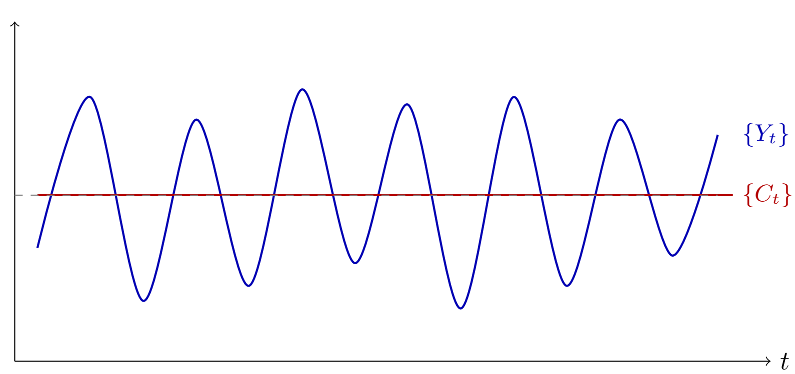

(consumption is constant if $\beta(1+r) = 1$)

The picture is the visual content of PIH: a wildly fluctuating income process can coexist with a perfectly smooth consumption process, because the household uses the saving/borrowing margin to absorb the income fluctuations.

10.2.4 Two Lessons of the PIH

The PIH theorem yields two policy-relevant implications, both deceptively simple.

Fact 10.2: The Timing of Income is Irrelevant

Two income streams $\{Y_t\}$ and $\{Y'_t\}$ that share the same present value PI generate the same consumption path $\{C_t\}$. Receiving income earlier vs. later, in lump sum vs. smoothly, in a recession vs. expansion—none of it matters for consumption, as long as PI is held fixed.

Fact 10.3: Ricardian Equivalence

Suppose the government finances a transfer (subsidy) of g dollars to households today, paid for by a tax of $(1 + r)^{\Delta T} g$ dollars in ΔT periods. Then PI is unchanged (the transfer's present value equals the tax's present value), so consumption is unchanged. The MPC out of such a fiscal transfer is exactly zero.

Remark (Why "Ricardian"?).

The principle was articulated (and rejected) by David Ricardo in 1820 and revived by Robert Barro in 1974. In a PIH world, debt-financed government spending is paid for by future taxes; rational households see this and save the entire transfer to pay the future tax. They are, in effect, internalizing the government's budget constraint into their own.

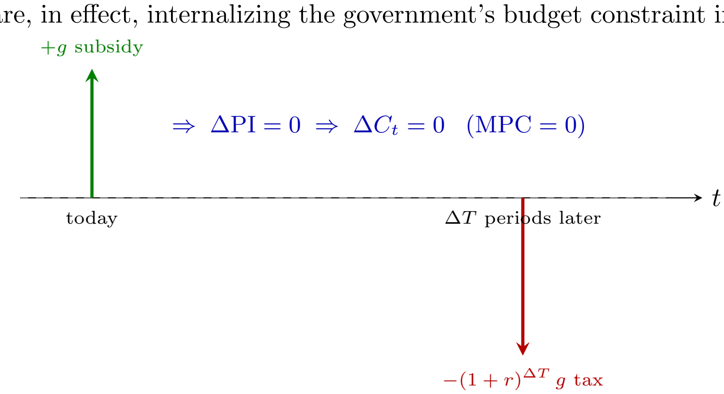

Ricardian Equivalence is one of the strongest predictions in macroeconomics, and its rejection in the data is one of the most important pieces of evidence *against* PIH in its purest form. Barro's (1974) revival made it operational by adding bequest motives, which extends the argument to overlapping generations. Empirically, MPCs out of transfers are typically estimated in the range 0.2–0.5, not 0, but the gap shrinks once liquidity constraints, finite lives, and uncertainty are layered on top of PIH.

10.2.5 Example 1: $\beta(1+r) = 1$ and Constant Consumption

The cleanest case parameterizes preferences and the interest rate so that the household weighs the present and the future symmetrically.

Setup

Assume $\beta(1+r) = 1$, equivalently $r = (1-\beta)/\beta$.

The Euler equation collapses:

$$u'(C_t) = \beta(1+r)\, u'(C_{t+1}) = u'(C_{t+1}) \implies C_t = C_{t+1} \qquad \forall t.$$

Consumption is **constant** across all periods. Plug $C_t = \bar{C}$ into the IBC:

$$\bar{C} \cdot \sum_{t=0}^{T} \frac{1}{(1+r)^t} = \text{PI} \implies \bar{C} = \phi(r) \cdot \text{PI},$$

where the *annuity factor* is

$$\boxed{\phi(r) \;=\; \frac{1 - \frac{1}{1+r}}{1 - \left(\frac{1}{1+r}\right)^{T+1}}.}$$

The household consumes the annuity value of its lifetime wealth in every period.

Remark (Limits of the Annuity Factor).

Two limits are worth noting:

- **Infinite horizon ($T \to \infty$):** $\phi(r) \to 1 - \frac{1}{1+r} = \frac{r}{1+r}$. The household consumes (roughly) the interest on its wealth and never depletes the principal.
- **Zero interest ($r \to 0$, finite T):** $\phi(r) \to \frac{1}{T+1}$. The household equally divides total wealth across the $T+1$ periods of life. (This is the Modigliani–Brumberg "life-cycle" intuition in its purest form.)

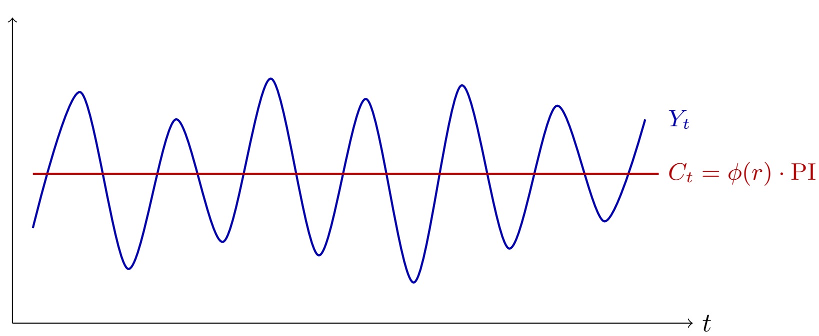

10.2.6 Example 2: General $\beta(1+r)$ and CRRA Preferences

When $\beta(1+r) \neq 1$, consumption is no longer constant. To get a clean closed-form, parameterize utility as constant relative risk aversion (CRRA):

$$u(C) = \begin{cases} \dfrac{C^{1-\sigma}}{1-\sigma} & \text{if } \sigma > 0,\ \sigma \neq 1, \\ \ln C & \text{if } \sigma = 1, \end{cases} \qquad \sigma > 0.$$

Then $u'(C) = C^{-\sigma}$. Plug into the Euler equation:

$$C_t^{-\sigma} = \beta(1+r)\, C_{t+1}^{-\sigma} \iff \left(\frac{C_{t+1}}{C_t}\right)^{\sigma} = \beta(1+r).$$

Take logs:

$$\boxed{g_C \;\equiv\; \ln\!\left(\frac{C_{t+1}}{C_t}\right) \;=\; \frac{1}{\sigma}\,\ln\!\big[\beta(1+r)\big].}$$

Consumption growth is constant and equal to $\frac{1}{\sigma}\ln[\beta(1+r)]$, regardless of t.

Definition 10.4: Intertemporal Elasticity of Substitution (IES)

The parameter $\frac{1}{\sigma}$ is the **intertemporal elasticity of substitution**: the percentage change in the consumption ratio C_{t+1}/C_t per percent change in the gross return $(1+r)$. CRRA preferences impose a constant IES (and a constant coefficient of relative risk aversion equal to σ); these two roles of σ are conflated under CRRA, which Epstein and Zin (1989) showed how to disentangle.

Fact 10.5: Lesson: Consumption Growth Reflects Preferences, not Income

The growth rate of consumption depends on:

- the patience of the household (β),
- the return on saving (r),
- the willingness to substitute consumption across time ($1/\sigma$).

It does **not** depend on the growth rate of income. Two households with the same preferences and same r will display the same consumption growth, even if one's income is rapidly rising and the other's is rapidly falling.

Remark (Three Cases of $\beta(1+r)$).

- $\beta(1+r) = 1$: $g_C = 0$, constant consumption (Example 1).
- $\beta(1+r) > 1$: market rewards saving more than the household discounts the future $\Rightarrow$ $g_C > 0$, consumption rises over time (the household is willing to start poor and end rich).

- $\beta(1+r) < 1$: the household discounts more than the market rewards $\Rightarrow g_C < 0$, consumption falls over time (the household front-loads consumption).

Empirically, U.S. aggregate consumption grows at roughly 2% per year, which combined with $r \approx 0.04$ requires β to be near (but slightly below) 1.

10.2.7 Connecting PIH to the Production Economy

The PIH framework uses an abstract asset A_t with a constant return r. In the production economies of Chapters 5–10, households save in a different form (capital k_t) which is rented out to a firm at rate r_t^k and depreciates at rate δ. These two pictures are equivalent up to a relabeling.

The household's period budget in the production setting is:

$$I_t + C_t \;=\; w_t L_t + r_t^k k_t,$$

where $I_t = k_{t+1} - (1-\delta)k_t$ is gross investment. Substituting:

$$k_{t+1} - (1-\delta)k_t + C_t \;=\; w_t L_t + r_t^k k_t,$$

i.e.,

$$\boxed{k_{t+1} + C_t \;=\; w_t L_t + \tilde{r}_t\, k_t, \qquad \tilde{r}_t \equiv 1 + r_t^k - \delta.}$$

With the identification $A_t \leftrightarrow k_t$, $Y_t \leftrightarrow w_t L_t$, and $1+r \leftrightarrow \tilde{r}_t$ (the gross return on capital, net of depreciation), this is exactly the asset-economy budget we have been using. Everything we have derived—IBC, Euler equation, PIH, Examples 1 and 2—applies verbatim to the production economy.

Remark (Why the Detour Through Production Matters).

Two reasons. First, the PIH framework as stated here treats r as exogenous, but in any general-equilibrium model with capital, r is endogenously determined by the marginal product of capital. The IBC is then a single household's budget; in equilibrium, r adjusts so that aggregate consumption and investment add up to aggregate output. Second, this connects PIH to RBC and Aiyagari: those models are GE versions of PIH plus production, plus (in Aiyagari's case) idiosyncratic income shocks. PIH provides the partial-equilibrium consumption function that the GE machinery then aggregates and equilibrates.

10.2.8 Empirical Application: Keynes (1936) and Friedman's Resolution

The PIH was not formulated in a vacuum. Friedman (1957) wrote it down explicitly to resolve a puzzle posed by Keynes (1936) and the empirical literature of the next two decades. Walking through the history in some detail does double duty: it shows the kind of empirical reasoning that the PIH framework enables, and it illustrates the broader methodological point that *looking at data through the lens of an economic model often changes what the* *data appear to say.*

Keynes (1936)

In *The General Theory*, Keynes wrote:

> "Consumption mainly depends on current income, and the relation is *fairly stable.*"

This is the so-called **Keynesian consumption function**: $C_i = a + bY_i$, with the marginal propensity to consume (MPC) b a structural parameter. Empirical support came from **cross-sectional** household surveys: in any given year, the projection of consumption on income across households produced a positive, statistically powerful, and seemingly stable slope $\hat{b}$ in the range 0.6–0.8. Richer households did consume more; poorer households consumed less; the relationship looked tight.

For roughly thirty years this was the dominant view of consumption behavior. The MPC of $\hat{b} \approx 0.75$ became a building block in textbook Keynesian multipliers, the IS-LM model, and policy analysis of fiscal stimulus.

The Puzzle

Two empirical facts began to undermine the Keynesian view by the 1940s and 1950s:

- In *long-run time series*, the aggregate saving rate S/Y was approximately *constant* at around 10% across the early 20th century, despite real per-capita income having grown several-fold. A fixed MPC of 0.75 would have predicted an ever-rising saving rate.
- In *short-run time series* (year-to-year aggregate fluctuations), consumption tracked income closely—but with a much smaller MPC than the cross-section implied.

The cross-section, the long run, and the short run gave *three different* answers for the same parameter b. The Keynesian theory, with a single structural MPC, could not accommodate all three.

Friedman's Resolution: Permanent vs. Transitory Income

Friedman's PIH provides the framework. Take Example 1 (Section 7) literally: each household i consumes a constant share of its lifetime wealth,

$$C_{it} = \phi(r)\,\mathrm{PI}_i, \qquad \forall t.$$

Now decompose observed income into a **permanent** and a **transitory** component:

$$Y_{it} \;=\; \underbrace{\phi(r)\,\mathrm{PI}_i}_{\equiv Y_{it}^P \text{ (permanent)}} \;+\; \underbrace{Y_{it}^T}_{\text{(transitory)}}.$$

Note that under PIH, C_{it} *is identical* to the permanent component: $C_{it} = Y_{it}^P$. The transitory component Y_{it}^T is what is left over (a windfall, a bonus, a temporary unemployment spell, an unexpectedly good harvest).

Computing Keynes's Coefficient Suppose we run the cross-sectional regression of C_i on Y_i at a fixed time t. The OLS slope is, by construction,

$$\hat{b} \;=\; \frac{\mathrm{Cov}[(,Y]_i\,,C_i)}{\mathrm{Var}[(]\,Y_i)}.$$

Substitute $Y_i = Y_i^P + Y_i^T$ and use bilinearity of covariance:

$$\hat{b} \;=\; \frac{\mathrm{Cov}[(,Y]_i^P\,,C_i) + \mathrm{Cov}[(,Y]_i^T\,,C_i)}{\mathrm{Var}[(]\,Y_i^P + Y_i^T)}.$$

Now invoke two structural assumptions implied by the PIH:

- Since $C_i = Y_i^P$, we have $\mathrm{Cov}[(,Y]_i^P\,,C_i) = \mathrm{Var}[(]\,Y_i^P)$.
- Assume **transitory income is independent of permanent income**: $Y_i^T \perp \mathrm{PI}_i$. Then $\mathrm{Cov}[(,Y]_i^T\,,C_i) = \mathrm{Cov}[(,Y]_i^T\,,Y_i^P) = 0$ and $\mathrm{Var}[(]\,Y_i^P + Y_i^T) = \mathrm{Var}[(]\,Y_i^P) + \mathrm{Var}[(]\,Y_i^T)$.

The Keynesian regression coefficient simplifies to

$$\boxed{\hat{b} \;=\; \frac{\mathrm{Var}[(]\,Y_i^P)}{\mathrm{Var}[(]\,Y_i^P) + \mathrm{Var}[(]\,Y_i^T)} \;=\; \text{share of permanent income in cross-sectional income variation.}} \qquad \text{(Friedman)}$$

Fact 10.6: $\hat{b}$ Is Not the MPC

The coefficient $\hat{b}$ that Keynes interpreted as the MPC is in fact a **signal-to-total ratio**: the share of cross-sectional income variation that comes from permanent differences across households. Under the PIH, the true MPC out of a transitory shock is essentially zero, and the MPC out of a permanent shock is $\phi(r) \approx r$ (small). Neither equals $\hat{b}$.

Why This Resolves the Puzzle

The Friedman decomposition explains all three empirical observations:

- **Cross-section:** Most of the variance *between households* at a point in time is permanent (one household is a doctor, another a janitor; this difference is not transitory). So $\mathrm{Var}[(]\,Y^P)$ is large relative to $\mathrm{Var}[(]\,Y^T)$, $\hat{b}$ is close to 1, and the cross-section delivers a large estimated slope. *Not because households consume out of current income, but because most of the income variation is permanent and fully passes through to consumption.*
- **Long-run time series:** The growth of aggregate income across decades is also permanent. A permanent rise in income of $X\%$ raises consumption by $X\%$, leaving S/Y constant. Consistent with the data.
- **Short-run time series:** Year-to-year aggregate fluctuations have a large transitory component (recessions, booms). PIH predicts that consumption barely responds to these—i.e., a small short-run MPC. Also consistent with the data.

Remark (Visualizing the Decomposition).

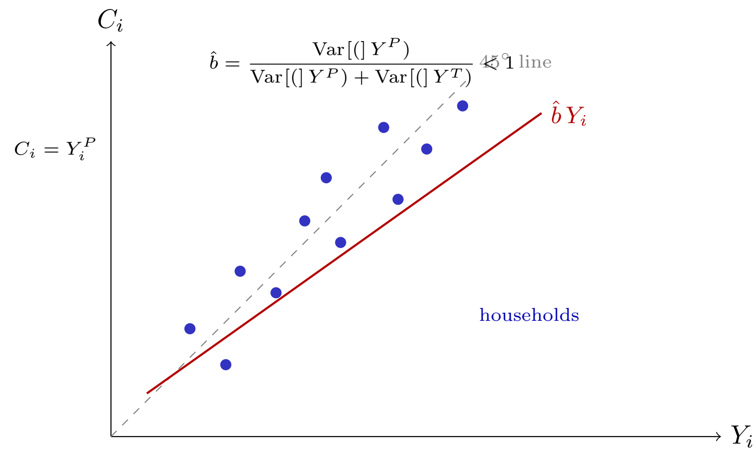

Each blue dot is a household. The vertical coordinate $C_i = Y_i^P$ is the permanent part of income; the horizontal coordinate $Y_i = Y_i^P + Y_i^T$ adds transitory noise. The OLS line through the cloud (red) has slope strictly less than 1 because the horizontal noise inflates the denominator $\text{Var}[(]\,Y_i)$ without inflating the numerator $\text{Cov}[(, Y]_i\,, C_i)$. The slope is exactly the signal-to-total ratio.

Remark (A Modern Coda: Measuring Permanent Income is Hard).

Identifying Y^P and Y^T in real data turns out to be one of the deepest empirical challenges in macro labor. With panel data on income (e.g., the PSID), one can decompose a worker's income time series into transitory and persistent components using statistical assumptions on the income process—but quantifying *subjective* permanent income (what the household believes about its lifetime wealth, taking into account future uncertainty) requires modeling the household's information set and forecasting behavior. Ludvig Straub (Clark Medal 2024) built much of his career on this kind of measurement, including a job-market paper that proposed a way to identify permanent-income shocks from observed updates in saving behavior. Three decades after Friedman, the issue is not closed.

10.3 Adding Uncertainty: Hall's Random Walk Hypothesis

The deterministic PIH is beautifully clean but obviously oversimplified: real income streams are stochastic, and a worker rarely knows even next year's income with certainty. We now make the income process **stochastic** and ask how the household optimizes when it must *forecast* its future. The headline, due to Hall (1978), is that under appropriate assumptions **consumption follows a random walk**: the best predictor of next period's consumption, conditional on everything known today, is today's consumption. This delivers a clean, testable empirical implication that organized two decades of subsequent work.

10.3.1 From PIH to Stochastic Income

What changes when we add uncertainty? Two things, both fundamental.

- **Forecasting becomes part of the problem.** The household no longer knows Y_{t+1}, only its conditional distribution given everything observed up to date t. The Euler equation acquires an expectation operator.
- **Precautionary saving emerges.** If u' is convex (i.e., $u''' > 0$), Jensen's inequality implies $\mathbb{E}_t[u'(C_{t+1})] > u'(\mathbb{E}_t[C_{t+1}])$, so a household facing risky future income saves more than one facing the deterministic-equivalent income. This is invisible in the deterministic PIH and shows up the moment uncertainty is added.

Hall (1978) is essentially the analog of PIH under uncertainty. Its central claim is simple: *in a rational, forward-looking household, the part of consumption growth that can be predicted from time-t information has been smoothed out already.* The only source of consumption changes is unexpected news.

10.3.2 Setup with Stochastic Income

The setup mirrors the deterministic case above but with one substantive change: the income process is now random.

- **Time:** infinite horizon, $t = 0, 1, 2, \ldots$
- **Income process:** an exogenous stochastic sequence $\{Y_t\}_{t=0}^{\infty}$. We do not yet specify the law of motion.
- **History:** let
$$s^t \equiv \{Y_\tau\}_{\tau=0}^{t}$$
denote the history of income realizations through date t. The household's information set at t is s^t.
- **Conditional expectation:** we write $\mathbb{E}_t[\cdot]$ as shorthand for $\mathbb{E}[\cdot \mid s^t]$, expectation conditional on the history known at t.
- **Asset:** a single risk-free bond, with constant gross return $1 + r$. Crucially: there are **no state-contingent claims**.

All choices the household makes at date t are functions of the history s^t, since that is all the household knows. We therefore write consumption and asset holdings as

$$C_t(s^t), \qquad A_{t+1}(s^t).$$

The Household's Problem

Maximize expected discounted lifetime utility:

$$\boxed{\max_{\{C_t(s^t),\, A_{t+1}(s^t)\}} \mathbb{E}_0\left[\sum_{t=0}^{\infty} \beta^t\, u\big(C_t(s^t)\big)\right]}$$

subject to the period budget constraint, which must hold *for every history*:

$$C_t(s^t) + \frac{A_{t+1}(s^t)}{1+r} \;=\; Y_t(s^t) + A_t(s^{t-1}), \qquad \forall\, t,\, \forall\, s^t,$$

non-negativity $C_t(s^t) \geq 0$, A_0 given, and the **no-Ponzi condition**:

$$\lim_{T\to\infty} \frac{A_{T+1}(s^T)}{(1+r)^T} \;\geq\; 0, \qquad \forall\, \{s^T\}_{T\to\infty}.$$

Remark (Why "$A_{t+1}(s^t)$" and not "$A_{t+1}(s^{t+1})$"?).

This subscript matters more than it looks. Writing A_{t+1} as a function of s^t encodes the assumption that the household must commit to its asset position *before* seeing tomorrow's income realization. Only one bond is available; its return is the same across all states tomorrow, and the household chooses a single quantity at t.

If we had instead allowed $A_{t+1}(s^{t+1})$—i.e., the asset position depends on tomorrow's realized state—we would have given the household a complete set of **state-contingent claims** (Arrow-Debreu securities). It could then guarantee any consumption pattern it wanted across all future states, and the problem would collapse back to the deterministic PIH from the previous section. The whole point of "uncertainty" biting is the absence of state-contingent claims; that is what the index s^t on A_{t+1} encodes.

This distinction is also the reason we need an Euler equation with an expectation operator (Section 4 below) rather than a Euler equation that holds state-by-state.

Remark (Reading the No-Ponzi Condition).

The no-Ponzi-game condition says: the present value of asset holdings cannot grow without bound, along any history. If it could, the household would be running a Ponzi scheme—borrowing forever and rolling debt at $1 + r$, paying it off by borrowing more—and would always strictly prefer to do that, since extra borrowing relaxes today's budget. Ruling this out is necessary for the problem to have a well-defined optimum.

In the deterministic finite-horizon PIH, the analog was simply $A_{T+1} \geq 0$. In the infinite-horizon stochastic case, no-Ponzi must hold along *every* possible history.

10.3.3 Optimality Conditions

Maximization with respect to $A_{t+1}(s^t)$, taking expectations over tomorrow's income realization, yields the **stochastic Euler equation**:

$$\boxed{u'\big(C_t\big) \;=\; \beta\,(1+r)\,\mathbb{E}_t\big[\,u'\big(C_{t+1}\big)\,\big] \qquad \forall\, t = 0, 1, 2, \ldots} \tag{Stoch. EE}$$

This is identical to the deterministic Euler equation *except* that $u'(C_{t+1})$ is now replaced by its conditional expectation. The household equates today's marginal utility (a known quantity) to the expected discounted marginal utility tomorrow.

The terminal condition is the **transversality condition**:

$$\boxed{\lim_{T\to\infty} \beta^T \, \mathbb{E}_0\big[\, u'\big(C_T\big)\, A_{T+1} \,\big] \;\leq\; 0.} \tag{TVC}$$

The TVC says: the discounted expected value of terminal asset holdings, weighted by marginal utility, must vanish (or be non-positive). It rules out the household being unboundedly wealthy in the limit, which would be wasteful from a utility-maximization perspective.

Remark (Sufficiency: Why TVC + Euler Solves the Problem).

> The Euler equation is a necessary first-order condition; it is satisfied by infinitely many candidate sequences $\{C_t(s^t)\}$, including suboptimal ones that, e.g., consume too little and accumulate unbounded assets. The TVC is what selects the unique optimum among them. Geometrically: the Euler equation pins down the *shape* of consumption (its expected growth rate), and the TVC pins down the *level* (no slack at infinity). Together with the budget constraints, they fully characterize the solution.

10.3.4 The Random Walk Hypothesis (General Form)

Hall's (1978) insight is to recognize the Euler equation as a martingale property. Define the **discounted marginal utility**:

$$X_t \;\equiv\; \beta^t \, u'(C_t).$$

Substitute into the stochastic Euler equation:

$$u'(C_t) = \beta\,(1+r)\,\mathbb{E}_t[\,u'(C_{t+1})\,] \iff \beta^t \, u'(C_t) = (1+r)\,\mathbb{E}_t[\,\beta^{t+1}\, u'(C_{t+1})\,].$$

With $\beta(1+r)$ on both sides, this can be rearranged depending on whether $\beta(1+r) = 1$. In the simplest case (which we will treat in Hall's specification below), $\beta(1+r) = 1$ and the equation reduces to

$$\boxed{X_t \;=\; \mathbb{E}_t[X_{t+1}].}$$

The Random Walk Hypothesis (General Form)

If $\beta(1+r) = 1$, the discounted marginal utility $\{X_t\} \equiv \{u'(C_t)\}$ is a **martingale**: its conditional expectation tomorrow equals its current realization. Equivalently,

$$\mathbb{E}_t\big[\,\Delta X_{t+1}\,\big] \;=\; \mathbb{E}_t\big[X_{t+1} - X_t\big] \;=\; 0.$$

Remark (The Economic Meaning: All Predictable Variation Has Been Smoothed).

> The martingale property is the formal statement of a deep economic intuition. A rational, forward-looking household equates expected marginal utilities across periods. So any *predictable* change in marginal utility—anything the household could see coming using time-t information—would represent an opportunity to do better by consuming more today (or tomorrow). At the optimum, no such opportunities are left:

any expected differences in marginal utility have been smoothed out.

What remains in $X_{t+1} - X_t$ is the **innovation**: the unexpected part, the news. The household optimizes *conditional on its information set*; what it could not have foreseen, it could not have insured against.

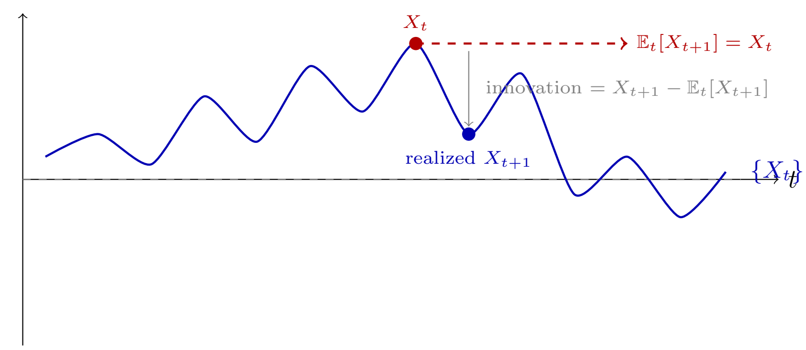

The realized path of X_t wanders (a true martingale), but at every point the conditional expectation of next period's value is exactly X_t itself.

10.3.5 Hall (1978): Quadratic Utility and the Random Walk in Consumption

So far we have shown that *marginal utility* is a martingale—an empirically inconvenient claim, since marginal utility is not directly observable. Hall's contribution is to add two assumptions that make the result testable in terms of the *level of consumption*, which is observable.

Hall's Assumptions

(1) Quadratic utility:

$$u(C) \;=\; C - \frac{a}{2}\,C^2, \qquad a > 0,$$

which gives marginal utility

$$u'(C) \;=\; 1 - a\,C, \qquad u''(C) = -a < 0, \qquad u'''(C) = 0.$$

(2) Patience matches the interest rate: $\beta(1+r) = 1$.

Plug into the stochastic Euler equation:

$$1 - a\,C_t \;=\; \mathbb{E}_t[\,1 - a\,C_{t+1}\,] \;=\; 1 - a\,\mathbb{E}_t[C_{t+1}].$$

The constants and the a cancel:

$$\boxed{C_t \;=\; \mathbb{E}_t[\,C_{t+1}\,].} \qquad \text{(Hall RWH)}$$

Hall's (1978) Random Walk Hypothesis

Under quadratic utility and $\beta(1+r) = 1$, the level of consumption itself is a martingale: $\{C_t\}_{t=0}^{\infty}$ is a **random walk**. The best predictor of next period's consumption is today's consumption.

Remark (Why Quadratic Utility?).

The general form of the random walk hypothesis says *marginal utility* is a martingale. To translate this into a statement about *consumption*, we need a relationship between the two:

$$u'(C_t) = \mathbb{E}_t[u'(C_{t+1})] \overset{?}{\Longrightarrow} C_t = \mathbb{E}_t[C_{t+1}].$$

The implication holds if and only if u' is *linear* in C. The unique strictly concave specification with linear u' is the quadratic. Under CRRA, log, or any other "standard" utility, u' is nonlinear, so the unobservable martingale in u' does *not* translate cleanly into a martingale in C. This is the strong assumption Hall is paying for testability.

A consequence of $u'(C) = 1 - a\,C$ is that $u''' = 0$: the third derivative vanishes. By a standard result, this implies **no precautionary saving motive**—the household is risk-averse but not prudent. Future income uncertainty changes the household's consumption distribution but not its level. This is one of the most criticized features of Hall's model and a major reason that later buffer-stock work (Carroll, 1997) adopted CRRA with $u''' > 0$.

Remark (The Tradeoff: Realism vs. Testability).

The lecturer's point in introducing Hall's specific assumptions is worth re-emphasizing: when you really want to test whether consumption behavior is rational, what you want to test is whether *marginal utility* is a martingale. That is the structural prediction. Translating this to a test on consumption levels requires a strong utility-function assumption that is itself unlikely to be exactly right. Empirical rejections of Hall's RWH might therefore be rejections of (i) rationality, (ii) $\beta(1 + r) = 1$, or (iii) quadratic utility—without further assumptions one cannot disentangle them.

10.3.6 Why It Matters: A Testable Implication

Hall's result is theoretically suggestive, but its real impact came from converting it into a sharp empirical test. If $C_t = \mathbb{E}_t[C_{t+1}]$, then the unexpected change $\varepsilon_{t+1} \equiv C_{t+1} - \mathbb{E}_t[C_{t+1}] = C_{t+1} - C_t$ should be **orthogonal** to all time-t information. In particular, lagged variables X_t in the household's information set should have no predictive power for ΔC_{t+1}:

$$\text{Cov}[(,\Delta]\,C_{t+1},\,X_t) \;=\; 0 \qquad \text{for all } X_t \text{ in the information set.}$$

This is the **Hall test**: regress ΔC_{t+1} on lagged income, lagged consumption growth, lagged stock returns, etc., and look for non-zero coefficients. Under PIH+RWH, all such coefficients should be zero. A statistically significant coefficient indicates **excess sensitivity**—consumption responding to predictable changes in income—which is the leading robust empirical violation of the model.

Remark (The Modern Empirical Frontier).

> The 1978–1990 literature on Hall's test (Flavin 1981, Campbell–Mankiw 1989, Zeldes 1989) found systematic excess sensitivity to predictable income changes, particularly for low-wealth households. Modern micro studies using natural experiments—tax rebates, lottery wins, oil-price shocks affecting region-specific incomes—continue to find MPCs out of transitory shocks in the range 0.2–0.5, much higher than the near-zero level predicted by PIH+RWH. The leading explanations involve liquidity constraints (Aiyagari-style buffer-stock saving), present bias (Laibson 1997), and cognitive limitations on forecasting permanent income. Section 10.4 below walks through the canonical results.

10.3.7 Quadratic Utility and Precautionary Saving

The previous remarks flagged $u''' = 0$ as "the strong assumption Hall is paying for testability." This subsection makes that flag operational: *quadratic utility itself has a testable implication*, namely the **absence of precautionary saving**. The literature has overwhelmingly rejected this implication, which is why modern consumption-saving work uses CRRA preferences and a *drifted* random walk in $\ln C_t$ rather than Hall's clean random walk in levels.

The Mechanism

Return to the stochastic Euler equation, written under $\beta(1+r) = 1$ for clarity:

$$u'(C_t) \;=\; \mathbb{E}_t[\, u'(C_{t+1})\,]\,.$$

Hold $\mathbb{E}_t[C_{t+1}]$ fixed and consider a **mean-preserving spread** of the conditional distribution of $C_{t+1} \mid s^t$ (intuitively, "more uncertainty about tomorrow"). What happens to today's optimal C_t? The answer depends entirely on the curvature of u':

- **If u' is linear in C (i.e., $u''' = 0$, the quadratic case).** Then $\mathbb{E}_t[u'(C_{t+1})] = u'(\mathbb{E}_t[C_{t+1}])$, so the right-hand side of the Euler equation is unchanged by the spread. Today's marginal utility, and therefore today's C_t, is also unchanged. *Pure mean-preserving uncertainty has no effect on optimal consumption today.*

- **If u' is convex in C (i.e., $u''' > 0$, the "prudent" case).** By Jensen's inequality, $\mathbb{E}_t[u'(C_{t+1})] > u'(\mathbb{E}_t[C_{t+1}])$. The mean-preserving spread strictly raises the right-hand side. To restore the equation, $u'(C_t)$ must rise; since $u'' < 0$, this means C_t falls. The household responds to higher uncertainty by **consuming less today and saving more**.

The second case is what we call **precautionary saving**: the household holds extra wealth today specifically to buffer against the variance of future consumption.

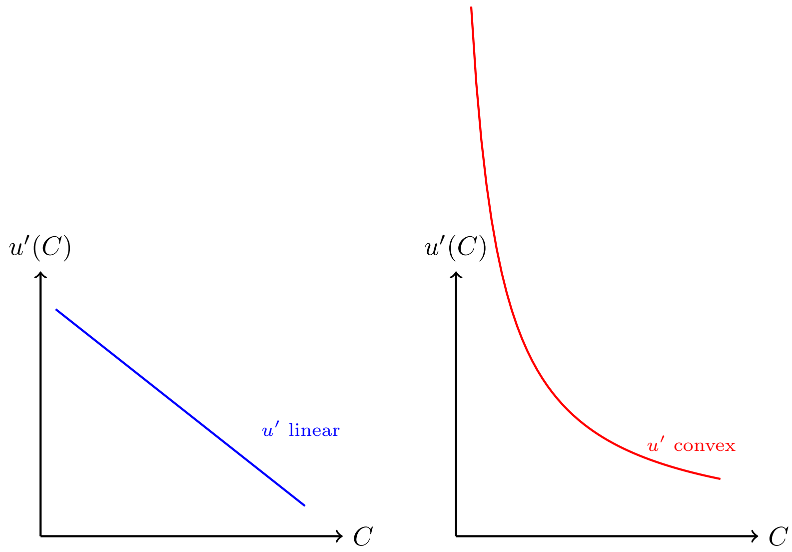

(a) Quadratic utility ($u''' = 0$)
No precautionary saving

(b) Prudent agent ($u''' > 0$)
Precautionary saving present

Remark (Precautionary Saving ≠ Risk Aversion).

A common confusion—surprisingly persistent even among trained economists—is to conflate precautionary saving with risk aversion. They are governed by different derivatives of u:

- **Risk aversion** is about $u'' < 0$: the household dislikes mean-preserving spreads in *utility levels*. A quadratic-utility household is risk-averse—welfare falls when income becomes more variable.

- **Prudence**, equivalently the precautionary-saving motive, is about $u''' > 0$: the household responds to mean-preserving spreads in *consumption* by *adjusting* how it smooths intertemporally. The Euler equation pins down expected marginal utility, not expected utility, so what governs the saving response is curvature one derivative deeper.

A quadratic-utility household is risk-averse but *not* prudent. It dislikes risk but does not save extra to buffer against it. CRRA, log, and exponential preferences are all both risk-averse and prudent.

Empirical Verdict and a Modern Replacement

The empirical consumption-saving literature has produced *robust* evidence that households increase their saving when their perception of future income risk rises—across many different proxies for risk (job-loss probability, occupation-level income variance, expected health shocks, even children's future labor-market prospects). For one influential example among many, **Boar (2017)** documents what she labels *dynastic precautionary saving*: parents who anticipate higher income risk for their children save more on their own behalf, with implications for overlapping-generations models of redistribution.

Remark (CRRA and the Drifted Random Walk).

Once $u''' > 0$ is allowed, C_t is no longer a martingale, but a closely related result survives. Under **CRRA preferences** ($u(C) = C^{1-\sigma}/(1-\sigma)$) and a strong assumption on the income-shock distribution (typically lognormal innovations to permanent income), one can derive that *log consumption follows a drifted random walk*:

$$\ln C_{t+1} \;=\; \ln C_t \;+\; \mu \;+\; \nu_{t+1}, \qquad \mathbb{E}_t \nu_{t+1} = 0,$$

where the drift μ depends on β, r, σ, and the variance of innovations. This is the modern formulation that replaces Hall's level random walk in quantitative work, and is the version derived in the homework. The testable implication—that anything in X_t should have zero predictive power for the residual ν_{t+1}—survives the transition; what changes is the functional form of the regression.

10.3.8 The Markov Simplification

The history-based formulation above is conceptually clean but operationally inconvenient: s^t grows without bound. In the chapters on RBC and BCA (and again in the Aiyagari chapter that follows) we used—and will use—the recursive formulation with a fixed-dimensional state instead. The bridge between the two is the **Markov assumption**.

Markov Assumption

Suppose the income process $\{Y_t\}$ is Markov: the conditional distribution of Y_{t+1} depends on s^t only through Y_t. Then for any two histories $\tilde{s}^t, \hat{s}^t$ with the same current realization $\tilde{Y}_t = \hat{Y}_t$ and the same current asset position $\tilde{A}_t = \hat{A}_t$, the household's optimal consumption-saving choice is identical.

In other words, the entire history s^t collapses to two state variables: current income Y_t and current assets A_t. The household's policy can be written as $C_t = g_C(A_t, Y_t)$, $A_{t+1} = g_A(A_t, Y_t)$. This is exactly the recursive structure used in the RBC model (where the state was (a_t, k_t)) and the Aiyagari model (where it was (z_t, k_t) for each household).

Remark (The History Form vs. the Recursive Form).

The history form $C_t(s^t)$ is more general: it works whether or not income is Markov. The recursive form $C_t = g_C(A_t, Y_t)$ requires the Markov assumption but is dramatically more tractable—it reduces an infinite-dimensional choice problem to a function on a fixed state space. All of the computational machinery (VFI, Aiyagari's outer-loop algorithm, etc.) requires the recursive form.

For empirical work the choice depends on the question: if you only need the Euler equation as a moment condition for testing, the history form is fine and imposes no Markov assumption. If you need to compute the policy function or the wealth distribution, you need Markov plus the recursive formulation.

10.4 Empirical Tests of PIH/RWH

The PIH and the random walk hypothesis are now in hand as theoretical predictions. We turn to the empirical literature that took them seriously as testable claims. The arc of this literature, spanning four decades, can be summarized as follows: aggregate tests rejected the representative-agent rational-expectations model; micro tests using natural experiments tightened the rejection by isolating clean variation in anticipated income; one major paper (Hsieh, 2003) appeared to rescue PIH but was later overturned by a more careful replication (Kueng, 2018). The cumulative verdict is that PIH/RWH systematically understates the consumption response to predictable income changes—the phenomenon now universally called **excess sensitivity**.

10.4.1 The Three Implications to Be Tested

Before reviewing the literature, it helps to organize the predictions into three sharp hypotheses about how consumption C_t should respond to different kinds of income changes ΔY_t:

Fact 10.7: Predictions of PIH/RWH

1. **Anticipated $\Delta Y_t \implies$ no ΔC_t.** If a change in income was already in the household's information set at time $t-1$, it has already been incorporated into Permanent Income, and therefore into C_{t-1}. The realized change today carries no news.

2. **Transitory unanticipated $\Delta Y_t \implies$ small ΔC_t.** An unexpected one-period shock raises Permanent Income only by the annuity value of the shock, which is small. Most of the windfall is saved (and consumed slowly over remaining life).

3. **Permanent unanticipated $\Delta Y_t \implies$ large ΔC_t.** A shock that revises the entire future income path one-for-one (e.g., a persistent productivity shift) raises Permanent Income by approximately the size of the shock itself, and consumption should respond nearly one-for-one.

The strongest and easiest-to-test prediction is (1): under PIH/RWH, consumption should *not* respond to news that was already public at $t-1$. This is the implication around which the entire empirical literature is organized.

In practice no shock is exactly "permanent." The third prediction is operationalized as: persistent unanticipated shocks (e.g., AR(1) innovations with ρ close to 1) generate a consumption response close to the size of the shock.

10.4.2 Aggregate Tests: Hall (1978) and Campbell–Mankiw (1989)

Hall's Original Test

Hall's (1978) original strategy is the obvious one: take aggregate U.S. consumption data, regress its first difference on lagged variables in the household's information set, and check whether anything predicts ΔC_{t+1}. Formally,

$$\Delta C_t \;=\; \alpha \;+\; X_t\,\beta \;+\; \varepsilon_t,$$

where X_t is a vector of variables known to the representative household at time t (e.g., lagged income Y_{t-1}, lagged consumption growth ΔC_{t-1}, lagged stock returns). The null hypothesis is RWH:

$$H_0: \quad \beta = 0.$$

Hall's finding was that for $X_t = (Y_{t-1}, C_{t-1})$, one *statistically* fails to reject $\beta = 0$. He interpreted this as evidence *for* PIH.

Remark (Statistical vs. Economic Significance).

> Failing to reject $\beta = 0$ is not the same as having shown $\beta = 0$. There are two ways the test can come out "negative":
>
> - **(Strong-evidence case.)** The point estimate $\hat{\beta}$ is genuinely close to zero *and* its standard error is small. The confidence interval excludes economically meaningful effects. This is strong evidence for PIH.
> - **(Underpowered case.)** The point estimate is small but the standard error is large. The confidence interval includes economically meaningful effects. This is *not* evidence for PIH; it is just an underpowered test.
>
> Hall's confidence intervals were narrow enough to make a serious argument, but already in his data the upper end of the CI was not literally zero. Subsequent papers focused on this gap.
>
> 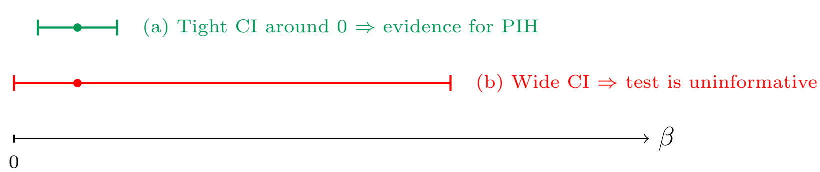
>
>
> Both panels have the same point estimate $\hat{\beta}$, but only (a) constitutes evidence for the null. Hall's and follow-up aggregate work landed somewhere between the two.

Campbell and Mankiw (1989): The Two-Type Model

Campbell and Mankiw (1989) sharpened the aggregate test by writing down a structural alternative to PIH. Suppose the population is split into two types:

- A fraction $\lambda \in (0, 1)$ of households is **hand-to-mouth**: they consume their income period-by-period, $C_t = Y_t$, so $\Delta C_t = \Delta Y_t$.
- The remaining fraction $1 - \lambda$ is fully **PIH**: they smooth, so ΔC_t is pure news, $\Delta C_t = \varepsilon_t$ with $\mathbb{E}_{t-1}\varepsilon_t = 0$.

Aggregating across the two types gives the regression

$$\boxed{\Delta C_t \;=\; \lambda\,\Delta Y_t \;+\; (1-\lambda)\,\varepsilon_t \;+\; X_t\,\gamma, \qquad \mathbb{E}_{t-1}\varepsilon_t = 0.} \tag{C\&M}$$

PIH is the special case $\lambda = 0$. A finding $\hat{\lambda} > 0$ rejects PIH and quantifies the share of consumption that responds to current income—i.e., excess sensitivity.

Remark (The Identification Problem).

> OLS estimation of (C&M) is biased: by construction ΔY_t is correlated with the PIH innovation ε_t (a shock to Y_t *is* the news that updates Permanent Income). Campbell and Mankiw therefore instrument ΔY_t with lagged variables Z_{t-1} that they argue are correlated with ΔY_t but uncorrelated with ε_t—typically lagged stock returns, lagged interest rates, lagged consumption growth. The strategy relies on strong time-series exogeneity assumptions, and the modern econometrics literature views the resulting instruments with considerable skepticism. This is one of the reasons the literature shifted to natural-experiment-based micro tests.

The headline result: $\hat{\lambda} \approx 0.5$ with a standard error of roughly 0.1–0.2. Half of aggregate consumption appears to be hand-to-mouth, decisively rejecting the representative-agent PIH at conventional levels.

Remark (What Aggregate Tests Can and Cannot Show).

> Two caveats are important. First, an aggregate $\hat{\lambda} \approx 0.5$ is consistent with multiple structural stories: a literal 50-50 split between two types, a more complex distribution with most households slightly excessively sensitive, a population with binding liquidity constraints in some periods, a representative agent with present-biased preferences, etc. The aggregate test rejects the *representative-agent rationalization* but does not pin down which structural model takes its place.
>
> Second, the natural correction is to move to micro data, where one can directly observe individual responses to anticipated income changes. The remainder of this section follows that route.

10.4.3 Excess Sensitivity in Micro Data

The micro literature focuses on the strongest PIH prediction: $\partial C_t / \partial(\text{anticipated } \Delta Y_t) = 0$. The empirical strategy is to find natural experiments in which households know about an upcoming income change and to compare consumption before and after the change actually arrives.

A representative sample of the literature:

- **Parker (1999):** Exploits the U.S. Social Security tax cap, which households hit deterministically partway through the year (after which take-home pay rises mechanically). Compares the consumption of households that hit the cap to those that do not. *Finds non-zero excess sensitivity.*
- **Souleles (1999):** Uses individual variation in the timing of federal tax refunds—refunds are a predictable, anticipated transfer. *Finds non-zero excess sensitivity.*
- **Browning and Collado (2001):** Uses Spanish data on civil servants who receive predictable end-of-year and summer bonuses. The regular, anticipated structure of the bonus

payments is precisely the type of variation PIH says should not move consumption. *Finds non-zero excess sensitivity.*

Across all three studies, the qualitative finding is the same: **ES** $\neq 0$. Households consume measurably more in the period when an anticipated income arrives, even though they could have smoothed perfectly. The micro literature, though heterogeneous in design, was nearly unanimous in this verdict.

10.4.4 Hsieh (2003): Alaska's Permanent Fund

Hsieh (2003) is the most cleanly identified test in the literature, and for a decade was treated as the definitive evidence *in favor of* PIH. The setup exploits Alaska's Permanent Fund Dividend (PFD), an annuity-style transfer paid to every Alaska resident.

Institutional Background

Alaska's Permanent Fund was created in 1976 as a sovereign-wealth-style fund: 25% of state oil-royalty revenue is deposited into the fund each year. Starting in 1982, 50% of the fund's annual dividend is paid out as the PFD, distributed equally to every individual Alaskan resident regardless of age. The total payout fluctuates year by year with oil-revenue and asset performance, and the per-person amount is determined by dividing the total by the resident count.

Information Timeline—The Source of Identification

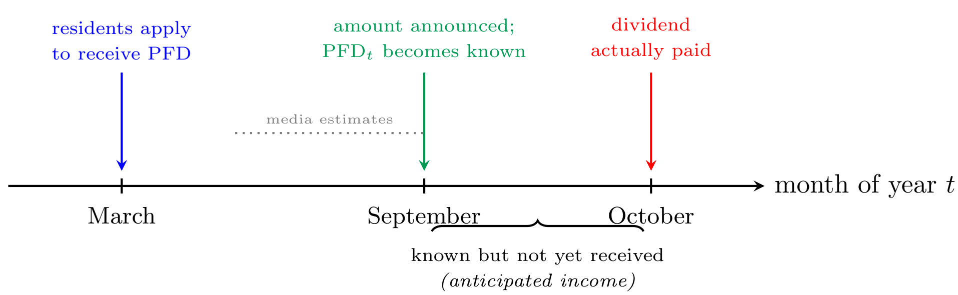

The key timing: by mid-September, every household knows exactly how much PFD it will receive (and rough media estimates have been circulating for months). The actual cash arrives in October. Under PIH, consumption should jump in September (or earlier, when the household first updates its expectations); it should *not* jump in October when an already-anticipated payment is mechanically deposited. The October consumption response is therefore the cleanest possible test of prediction (1).

Specification

For household h in year t, regress the log-change in consumption between Q3 and Q4:

$$\boxed{\ln\left(\frac{C_{t,h}^{Q4}}{C_{t,h}^{Q3}}\right) \approx \alpha \cdot \frac{\text{PFD}_t \times \text{family size}_h}{\text{family income}_{t,h}} + X_{t,h}\beta + \varepsilon_{t,h}.} \tag{10.1}$$

The right-hand-side regressor scales the dollar amount of dividend (per-person PFD times family size) by family income, giving the proportional income shock. Three sources of identifying variation feed into this exercise:

- *Across years*: PFD_t varies with oil prices and fund returns.
- *Across households*: family size differs, so the dollar amount differs.
- *Across households*: family income normalizes the shock into a proportional change.

This is effectively a difference-in-differences design with continuous treatment intensity. PIH predicts $\alpha = 0$.

Findings

Spending category	$\hat{\alpha}$	SE
Non-durables	0.0002	0.0324
Durables	0.166	0.088

For non-durables—the category closest to the theoretical C_t—the point estimate is essentially zero, and even three standard errors give an upper bound below 10%. Hsieh interpreted this as a clean failure to reject PIH. For durables, the estimate is meaningfully positive ($\hat{\alpha} \approx 0.17$), but Hsieh argued this is consistent with theory: durables are an investment good, and their bunched purchase in the dividend month is an optimal lumpy response, not a violation of consumption smoothing in the flow-utility sense.

Remark (Why This Was a Big Deal).

The earlier literature (Parker, Souleles, Browning–Collado) had found excess sensitivity, but each study could be challenged on identification grounds—small samples, weak instruments, idiosyncratic features of the institutional setting. Hsieh's design was unusually clean: large dollar amounts (substantial fraction of household income for many Alaskans), perfect anticipation by October, transparent variation across households, and a sharp prediction. For a decade, the verdict was that the previous excess-sensitivity findings were artifacts and PIH actually held in well-measured data. David Romer's textbook cited Hsieh as state-of-the-art evidence for PIH.

10.4.5 Kueng (2018) Revisits Hsieh

The story did not end there. Kueng (2018) replicated Hsieh's specification with three modifications, each motivated by a methodological concern:

1. **Normalize by expenditure rather than income.** Income data have well-known measurement problems (under-reporting, missing saving choices, transitory components). Total expenditure is widely viewed as a more reliable proxy for permanent income in micro data, so the proportional shock is rescaled by total household expenditure.

2. **Extend the sample by 12 years.** Hsieh used the original CEX panel covering only the early Alaska Permanent Fund years; Kueng adds more than a decade of additional data, sharply increasing statistical power.

3. **Use non-Alaskans as a control group.** Hsieh's identifying variation came entirely *within* Alaska—across family sizes and incomes. This is potentially problematic if Alaskans have unobserved characteristics (oil-economy seasonality, idiosyncratic Q4 spending patterns) that contaminate the estimate. Kueng adds non-Alaskan households as a control group: people who, by construction, do not receive the PFD and therefore should show no Q3-to-Q4 jump driven by it.

The modification (3) is the most consequential. With non-Alaskans absorbing common Q3-to-Q4 spending patterns and Alaskans absorbing the same patterns plus the PFD response, the difference-in-differences identifies the true PFD effect cleanly.

Findings

After all three modifications, Kueng finds **statistically and economically significant excess sensitivity** for non-durables in Alaska—reversing Hsieh's central conclusion. The PFD, despite being perfectly anticipated by October, generates a measurable consumption response when paid out.

Remark (Identification: A Cautionary Note).

> Kueng's identification is itself imperfect. Comparing Alaskans to non-Alaskans introduces a new concern: any difference between the two groups beyond the PFD could be loaded onto $\hat{\alpha}$. Kueng works hard to control for observable characteristics, but as a matter of principle the tradeoff is between Hsieh's tighter sample (no non-Alaskan contamination, but potentially under-identified) and Kueng's broader sample (cleaner counterfactual, but reliance on observables to align the two groups). This is the natural state of empirical macro—no single design is perfect, and what carries the field forward is whether successive studies converge in spite of their individual identification weaknesses.

10.4.6 Where the Literature Stands

The four-decade arc tells a consistent story once we put all the evidence together:

- Aggregate tests (Hall, Campbell–Mankiw) reject the representative-agent PIH and quantify excess sensitivity at $\hat{\lambda} \approx 0.5$.
- Pre-Hsieh micro tests (Parker, Souleles, Browning–Collado) find non-zero excess sensitivity in cleanly identified natural experiments.
- Hsieh (2003) appeared to overturn this verdict using the Alaska PFD design.
- Kueng (2018) re-establishes excess sensitivity by adding a control group, extending the sample, and using a more reliable normalization.

The cumulative finding is that anticipated income changes *do* move consumption, in violation of PIH. The leading structural explanations—each motivating a research program in its own right—are:

- **Liquidity constraints.** Households cannot freely borrow against future income. When a transfer arrives, constrained households mechanically increase spending. This is the channel emphasized by buffer-stock-saving models (Carroll, 1997) and by Aiyagari-style heterogeneous-agent macro models (next chapter).
- **Present bias.** Hyperbolic discounting (Laibson, 1997) generates excess sensitivity in a representative-agent setting without market frictions.
- **Inattention and information frictions.** Households may not fully process anticipated changes until the cash arrives, blunting forward-looking smoothing.
- **Mental accounting.** Behavioral consumers may treat "windfalls" (tax refunds, dividends) as a separate budget category from regular income.

Remark (The Methodological Lesson).

Read this less as a settled empirical matter and more as a model of how macro research progresses. Hsieh wrote a beautiful paper that overturned a literature; Kueng wrote an equally beautiful paper that overturned Hsieh; both will likely be overturned in turn. What matters is not that any one estimate is the truth, but that successive iterations of theory and measurement push the field toward a useful approximation. The Aiyagari and HANK frameworks of subsequent chapters can be read as theoretical responses to the empirical regularity, distilled across these papers, that consumption is excessively sensitive to anticipated income.

Remark (Chapter Summary).

- **The PIH theorem.** With perfect intertemporal financial markets, optimal consumption depends on income only through the present value of lifetime resources (Permanent Income, PI). The timing of income, holding PI fixed, is irrelevant.
- **Hall's RWH (1978).** Under quadratic utility and $\beta(1+r) = 1$, C_t is a martingale: any predictable change in consumption from time-t information has been smoothed out already. Innovations come only from news about PI.
- **Quadratic utility is testable, and it fails.** With $u''' = 0$ there is no precautionary saving; with $u''' > 0$ households cut consumption and save in response to mean-preserving uncertainty—a robust empirical regularity. CRRA + lognormal innovations gives the modern drifted-random-walk formulation.
- **The empirical literature systematically rejects PIH/RWH.** Aggregate tests yield $\hat{\lambda} \approx 0.5$ hand-to-mouth share; micro tests (Parker, Souleles, Browning–Collado) find non-zero excess sensitivity to anticipated income changes. Hsieh (2003) appeared to rescue PIH using the Alaska PFD, but Kueng (2018) overturned the result with a cleaner control group.
- **Structural explanations.** Liquidity constraints (Aiyagari-style buffer-stock saving), present bias (Laibson), inattention, and mental accounting all generate excess sensitivity. The Aiyagari and HANK models in the next chapter quantify these channels.

Chapter 11

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 **Aiyagari (1994)** 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.

11.1 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 11.1: Aggregate vs. Idiosyncratic Risk: A New Perspective on the Equity Premium

Mehra and Prescott (1985) 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 11.2: 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.

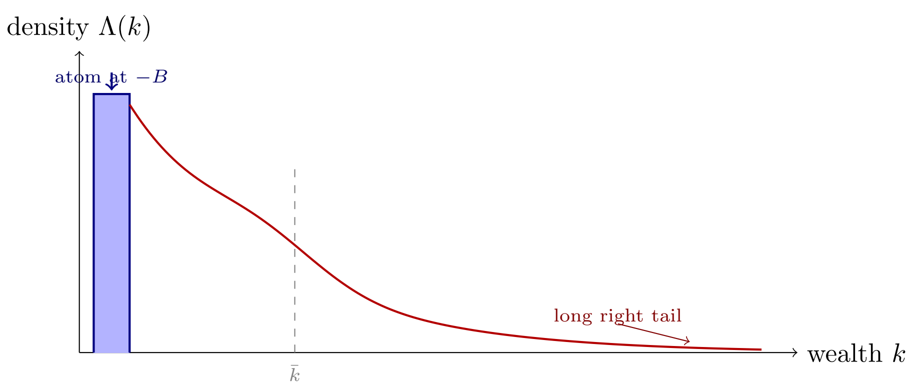

The atom at the borrowing constraint $-B$ is a generic feature of the model: a positive mass of households runs out of savings, hits the constraint, and stays there until a positive income shock pushes them off. We will see this atom emerge naturally from the algorithm.

11.2 The Environment

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

- **Time:** discrete, $t = 0, 1, 2, \ldots$
- **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 $\mathbb{Z} = \{z_1, \ldots, z_{nz}\}$.
- **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.

11.3 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.

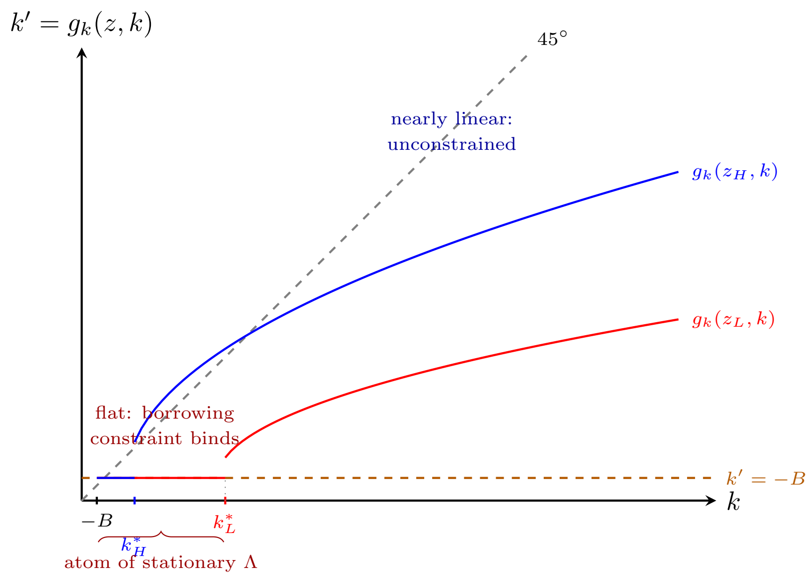

> - **Flat segment at $k' = -B$.** For each income level z, there is a threshold k_z^* below which the household would optimally choose $k' < -B$ in the absence of the constraint. The constraint forces these households to set $k' = -B$ exactly. Households who land in the flat region therefore *accumulate 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).

11.4 Stationary Recursive Competitive Equilibrium

We can now define the equilibrium concept formally.

Definition 11.3: 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:

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

11.5 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.

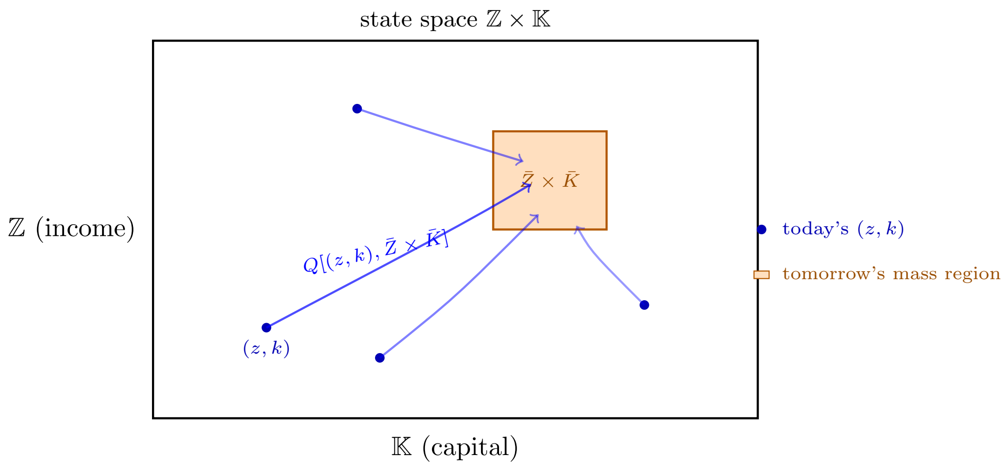

The diagram is the visual content of the stationarity equation: integrate the inflow probability Q over all today-states (blue dots), get tomorrow's mass in the orange box. Stationarity demands this equal $\Lambda(\bar{Z}\times\bar{K})$ itself.

11.6 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:

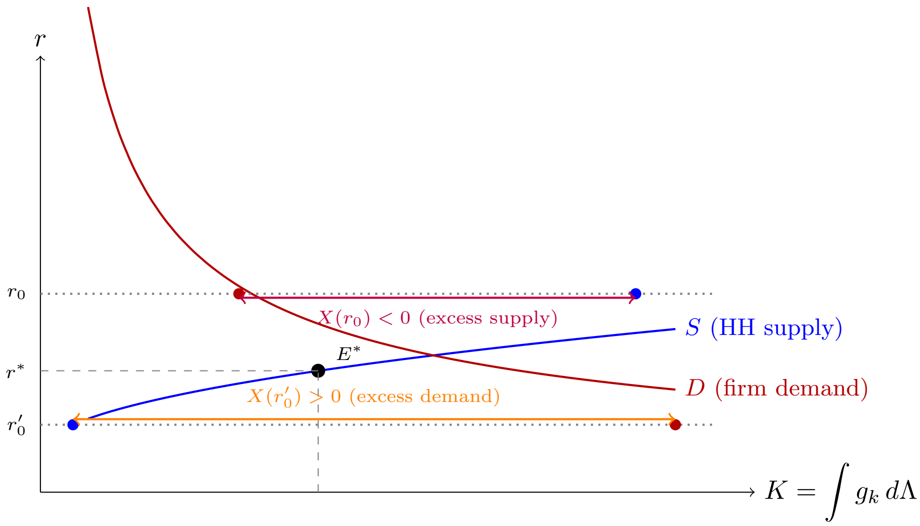

The picture clarifies the bisection logic: if at r_0 excess demand is negative (households want to save more than firms want to use), lower r_0; if it is positive, raise it.

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 $r^{\text{RA}} = 1/\beta - 1$ (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.

11.7 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:

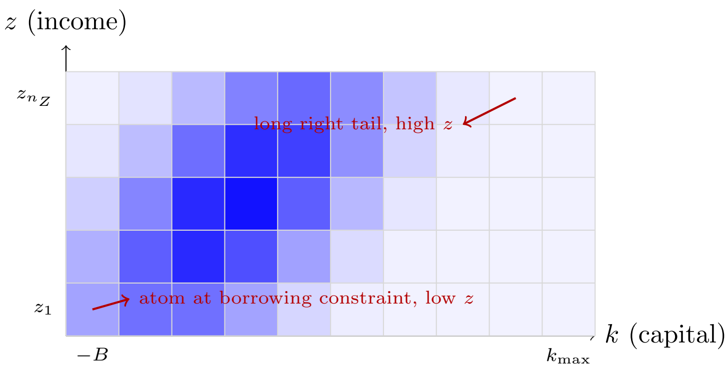

Mass concentrates along a positive ridge from low-z/low-k to high-z/high-k: persistently lucky households accumulate capital, persistently unlucky ones dissave to the borrowing limit. The marginal over k (summing across rows) is the wealth distribution figure from §1.

11.8 Visualizing the Algorithm

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

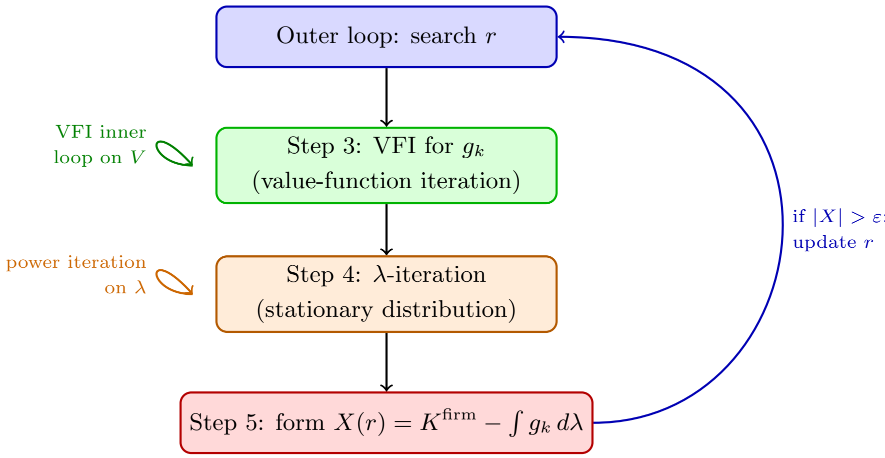

Three loops, three convergence criteria. The outer loop terminates when excess demand for capital is small. The VFI loop terminates when the value function stops moving. The λ-loop terminates when the distribution stops moving. Each loop sees the result of the inner loops as a fixed input.

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, Carroll 2006).** 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.

11.9 Beyond the Basic Aiyagari Model

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

- **Aiyagari (1994) 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: Krusell and Smith (1998).** 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.** Achdou et al. (2022) 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.

Chapter 12

Final Exam Review Guide

Remark (How to Use This Chapter).

This chapter is a compressed reference for the final exam. The exam covers two kinds of material in roughly equal proportion:

- **Theory drill** (Section 12.1): Six question types modeled on Problem Set 5 (chapter "Problem Set 8" in this book). Each subsection summarizes the lecture content needed to solve one type of question, with the canonical derivation worked through.
- **Computation drill** (Section 12.2): Three numerical algorithms from Part II—deterministic Value Function Iteration, stochastic VFI for the RBC model, and the Aiyagari outer loop. The exam asks for the algorithm in writing (steps and rationale), *not* executable code.

The Quick Reference Cheat Sheet in Section 12.3 collects the formulas you should be able to reproduce from memory. The earlier chapters of Part II contain everything in this review and more; this chapter prioritizes only what you need to solve PS5-style problems and write down computational procedures.

12.1 Theory Drill

Each subsection below is organized as **Question** (what an exam-style problem asks), **Setup** (the full model with all primitives stated), **Solution** (the canonical derivation), and **Punchline** (what to remember). The structure mirrors the six problems on Problem Set 5; if you can reproduce these derivations from scratch, you can solve the corresponding exam questions.

12.1.1 Solow Model with Population and TFP Growth

Question. The basic Solow model in the lectures has constant TFP and a constant labor force, and predicts *no* long-run growth. What if we let both TFP and population grow at constant exponential rates? Is there still a balanced growth path (BGP) where capital per worker grows at a constant rate? If so, how does that rate depend on the model's primitives?

Setup. The standard Solow ingredients:

- Cobb–Douglas production: $Y_t = A_t K_t^\alpha L_t^{1-\alpha}$, with $\alpha \in (0,1)$.
- Constant saving rate $s \in (0,1)$: investment $I_t = sY_t$.
- Capital accumulation: $K_{t+1} = I_t + (1-\delta)K_t$, with $\delta \in (0,1)$.
- Two new ingredients: TFP grows at rate g_A, i.e., $A_t = A_0 e^{g_A t}$; labor grows at rate g_L, i.e., $L_t = L_0 e^{g_L t}$.

You are asked to (i) show that a BGP for capital exists, and (ii) compute the BGP growth rates of total capital g_K and capital per worker g_k.

Solution. The technique is **guess and verify**: conjecture that on the BGP, $k_t \equiv K_t/L_t$ grows at some constant rate g_k, then plug the conjecture into the law of motion and demand consistency.

Step 1: Per-worker law of motion. Divide $K_{t+1} = sA_t K_t^\alpha L_t^{1-\alpha} + (1-\delta)K_t$ by L_{t+1} and use $L_{t+1}/L_t = e^{g_L}$:

$$e^{g_L} k_{t+1} \;=\; sA_t\, k_t^\alpha + (1-\delta)\, k_t. \tag{$*$}$$

Step 2: Substitute the BGP guess. Plug $k_t = k_0 e^{g_k t}$ and $A_t = A_0 e^{g_A t}$ into $(*)$, divide both sides by $k_0 e^{g_k t}$:

$$e^{g_L + g_k} \;=\; sA_0\, k_0^{\alpha-1}\, e^{(g_A + (\alpha-1)g_k)\,t} + (1-\delta).$$

Step 3: Demand time-invariance. The left side is constant in t. For the right side to also be constant, the coefficient of t in the exponent must vanish:

$$g_A + (\alpha-1)\,g_k = 0 \implies \boxed{g_k = \frac{g_A}{1-\alpha}, \qquad g_K = g_L + g_k = g_L + \frac{g_A}{1-\alpha}.}$$

Punchline.

- TFP growth is amplified into per-worker capital growth by the multiplier $1/(1-\alpha) > 1$. The mechanism: a TFP gain raises the marginal product of capital, calling forth additional investment, which raises the marginal product further—a within-period feedback that the Cobb–Douglas exponent α controls.
- Population growth contributes one-for-one to *total* capital ($g_K = g_L + \cdots$) but not at all to capital *per worker*. "Spreading more capital across more workers" offsets the additional accumulation, leaving k_t unaffected.
- The technique (substitute exponential guess, force the t-exponent to zero) is the standard tool for any BGP question. Memorize it.

12.1.2 Two-Period Saving with Risky Returns

Question. A consumer lives for two periods with log utility. They earn labor income in either period 1 or period 2 (two distinct sub-questions), and the gross return $1+r$ they face on saving/borrowing may be either certain or random. The exam asks: does the risk in r

matter for first-period consumption C_1? In particular, if r becomes uncertain while keeping $\mathbb{E}(r)$ fixed (a mean-preserving spread), does C_1 rise, fall, or stay the same?

The economic content: this is a question about precautionary saving versus the income/substitution-effect cancellation. Whether C_1 moves depends on a delicate interplay between (i) when income arrives and (ii) the curvature of the utility function in the random variable r.

Setup. Lifetime utility $U = \ln C_1 + \mathbb{E}[\ln C_2]$. Two scenarios:

Case A: income upfront. Endowment $Y_1 > 0$ in period 1 and zero in period 2. The agent saves $Y_1 - C_1$ at gross return $1 + r$, so $C_2 = (1 + r)(Y_1 - C_1)$.

Case B: income back-loaded. Endowment 0 in period 1 and $Y_2 > 0$ in period 2. The agent borrows C_1 in period 1 and repays $(1+r)C_1$ in period 2, so $C_2 = Y_2 - (1+r)C_1$.

In both cases, derive (i) the FOC for C_1, and (ii) the response of C_1 to a mean-preserving spread in r.

Solution: Case A (saver).

Lifetime utility:

$$U \;=\; \ln C_1 + \mathbb{E}[\ln(1 + r)(Y_1 - C_1)] \;=\; \ln C_1 + \ln(Y_1 - C_1) + \mathbb{E}[\ln(1 + r)].$$

The crucial observation: log utility makes the random variable r enter *additively* through $\mathbb{E}[\ln(1 + r)]$, which does not depend on C_1. The FOC therefore drops r entirely:

$$\frac{1}{C_1} \;=\; \frac{1}{Y_1 - C_1} \;\Longrightarrow\; C_1 = \tfrac{1}{2}Y_1.$$

A mean-preserving spread in r changes only $\mathbb{E}[\ln(1 + r)]$, which the FOC ignores. Hence C_1 is *unchanged.*

Solution: Case B (borrower).

Lifetime utility:

$$U \;=\; \ln C_1 + \mathbb{E}\big[\ln(Y_2 - (1 + r)C_1)\big].$$

Now r multiplies C_1 inside $\ln(\cdot)$, so r enters the FOC:

$$\frac{1}{C_1} \;=\; \mathbb{E}\left[\frac{1 + r}{Y_2 - (1 + r)C_1}\right] \;\equiv\; \mathbb{E}[f(r)].$$

Compute the curvature of f. By the quotient rule:

$$f'(r) = \frac{Y_2}{[Y_2 - (1 + r)C_1]^2}, \qquad f''(r) = \frac{2Y_2C_1}{[Y_2 - (1 + r)C_1]^3} > 0,$$

provided $C_2 > 0$ in every state. By Jensen's inequality, an MPS in r at constant $\mathbb{E}(r)$ strictly raises $\mathbb{E}[f(r)]$. The FOC then forces $1/C_1$ to rise—i.e., C_1 *falls.*

Punchline.

- In **Case A**, log utility and saving combine to make C_1 independent of return uncertainty. The substitution effect (higher expected return $\Rightarrow$ save more) and income effect (higher lifetime wealth $\Rightarrow$ consume more) cancel exactly under log preferences. *No precautionary response.*

- In **Case B**, the agent is borrowing, and uncertainty in r makes the marginal cost of borrowing $f(r) = (1+r)/(Y_2 - (1+r)C_1)$ stochastic. Convexity of f in r means the bad states (high r, very low C_2, very high marginal utility) outweigh the good states. *Borrowing falls; C_1 strictly decreases.*

- **General lesson for the exam.** The direction of the response to MPS in r depends on whether r enters the FOC additively (Case A; no effect on C_1) or multiplicatively (Case B; effect via Jensen on a convex function).

12.1.3 RBC: Bellman, Euler, and Linear-Quadratic Guess-and-Verify

Question. This subsection answers two questions a typical RBC problem combines.

(1) **General theory.** Given a stochastic neoclassical growth model with state (a, k), write down the Bellman equation, derive the Euler equation, and (under quadratic utility) show that consumption follows Hall's random-walk hypothesis.

(2) **Closed-form solution.** For a particular linear-quadratic specification (Blanchard–Fischer), guess that consumption is linear in capital and the productivity shock, and solve for the coefficients of the policy function.

Setup for (1): General stochastic Bellman. A planner faces:

- Per-period utility $u(c)$ (smooth, concave, satisfying Inada conditions).
- Cobb–Douglas production with stochastic TFP: $y_t = e^{a_t} f(k_t)$ where a_t follows an AR(1).
- Standard capital accumulation: $k_{t+1} = e^{a_t} f(k_t) + (1-\delta)k_t - c_t$.
- Discount factor $\beta \in (0, 1)$.

Solution for (1): Bellman, FOC, Euler.

$$V(a,k) \;=\; \max_{c,k'} \big\{u(c) + \beta\, \mathbb{E}[V(a',k') \mid a]\big\}, \quad k' = e^a f(k) + (1-\delta)k - c.$$

FOC w.r.t. k'. The first-order condition (using $\partial k'/\partial c = -1$):

$$u'(c) \;=\; \beta\, \mathbb{E}[V_{k'}(a',k') \mid a].$$

Envelope theorem. Differentiating the Bellman in k:

$$V_k(a,k) \;=\; u'(c) \cdot \big[e^a f'(k) + 1 - \delta\big].$$

Combining. Forward-shift the envelope ($V_{k'}(a',k') = u'(c') \cdot [e^{a'} f'(k') + 1 - \delta]$) and substitute into the FOC:

$$\boxed{u'(c_t) \;=\; \beta\, \mathbb{E}_t\big[u'(c_{t+1})\,(e^{a_{t+1}} f'(k_{t+1}) + 1 - \delta)\big].}$$

Hall's RWH (special case). If $u(c) = c - \theta c^2$ (so $u'(c) = 1 - 2\theta c$ is linear in c) and the gross return is constant at $1 + r$ with $\beta(1 + r) = 1$, the Euler equation collapses:

$$1 - 2\theta c_t \;=\; \mathbb{E}_t[1 - 2\theta c_{t+1}] \iff c_t = \mathbb{E}_t[c_{t+1}].$$

Consumption is a martingale; changes from c_t to c_{t+1} are unforecastable from time-t information.

Setup for (2): Linear-quadratic RBC (Blanchard–Fischer). A specific specification where the Euler simplifies:

- Period utility $u(C) = C - \theta C^2$, $\theta > 0$.
- Discount: the household discounts at rate $\rho > 0$, so the discount factor is $1/(1+\rho)$.
- Linear production with additive shock: $Y_t = AK_t + e_t$, with $A = \rho$ (a parameter restriction that simplifies the algebra).
- No depreciation: $K_{t+1} = K_t + Y_t - C_t = (1+\rho)K_t + e_t - C_t$.
- AR(1) shock: $e_t = \phi e_{t-1} + \varepsilon_t$, with $|\phi| < 1$ and ε_t i.i.d. mean-zero.

Question: guess $C_t = \alpha + \beta K_t + \gamma e_t$, and solve for (α, β, γ).

Solution for (2): Guess and verify. The Euler equation in this LQ environment simplifies to $C_t = \mathbb{E}_t[C_{t+1}]$ (Hall's RWH from above, with $\beta(1+r) = 1$ effectively imposed by $A = \rho$).

Step 1: Substitute the linear policy guess into the resource constraint. If $C_t = \alpha + \beta K_t + \gamma e_t$,

$$K_{t+1} \;=\; (1+\rho)K_t + e_t - (\alpha + \beta K_t + \gamma e_t) \;=\; (1+\rho-\beta)K_t + (1-\gamma)e_t - \alpha.$$

Step 2: Substitute into the Euler $C_t = \mathbb{E}_t[C_{t+1}]$.

$$\alpha + \beta K_t + \gamma e_t \;=\; \mathbb{E}_t[\alpha + \beta K_{t+1} + \gamma e_{t+1}] \;=\; \alpha + \beta\,\mathbb{E}_t[K_{t+1}] + \gamma\phi e_t.$$

Plug in $\mathbb{E}_t[K_{t+1}] = (1+\rho-\beta)K_t + (1-\gamma)e_t - \alpha$:

$$\alpha + \beta K_t + \gamma e_t \;=\; \alpha + \beta\big[(1+\rho-\beta)K_t + (1-\gamma)e_t - \alpha\big] + \gamma\phi e_t.$$

Step 3: Match coefficients.

- *K_t:* $\beta = \beta(1+\rho-\beta) \Longrightarrow \beta = \rho$ (the non-trivial root).
- *e_t:* $\gamma = \beta(1-\gamma) + \gamma\phi \Longrightarrow \gamma(1+\beta-\phi) = \beta \Longrightarrow \gamma = \rho/(1+\rho-\phi)$.
- *Constant:* $\alpha = \alpha - \beta\alpha \Longrightarrow \alpha = 0$ (since $\beta = \rho > 0$).

Result: $C_t = \rho K_t + \frac{\rho}{1+\rho-\phi}\, e_t$.

Punchline.

- The general FOC and envelope theorem deliver the Euler equation in any neoclassical model. Memorize the procedure: FOC w.r.t. k', then envelope theorem, then forward-shift.
- Hall's random-walk result requires both quadratic utility and $\beta(1+r) = 1$. Quadratic alone makes u' linear; $\beta(1+r) = 1$ kills the $\beta(1+r)$ factor in the Euler equation.
- For LQ-RBC, the policy function is linear in the state. The four-step procedure (guess $\to$ resource $\to$ Euler $\to$ match coefficients) is exam-mechanical—no creativity required, only careful algebra.

12.1.4 PIH and the Farmer/Non-Farmer Empirical Puzzle

Question. The empirical observation: in cross-sectional household data, farmers' average income is lower than non-farmers', and farmer income fluctuates more from year to year. If you ran a regression of consumption C_i on current income Y_i separately for the two groups, the regression slopes would differ. The exam asks: under the permanent-income hypothesis (PIH), what does theory predict about how the two slopes differ, and how do we explain the prediction?

The framing matters for the exam: the question is about *the apparent MPC*, the slope $\hat{b}$ in a regression of C_i on Y_i. PIH does *not* say that the apparent MPC is constant across groups—in fact, it predicts the apparent MPC depends on the income process.

Setup. Decompose individual income into permanent and transitory components:

$$Y_i \;=\; Y_i^P + Y_i^T, \qquad \mathrm{Cov}[(,Y]_i^P, Y_i^T) = 0.$$

Under the PIH, consumption is a constant fraction of permanent income:

$$C_i \;=\; \phi\, Y_i^P, \qquad \phi = r/(1+r) \text{ in the infinite-horizon case.}$$

You are asked to compute the cross-sectional regression slope $\hat{b} = \mathrm{Cov}[(,Y]_i, C_i)/\,\mathrm{Var}[(]\,Y_i)$ and interpret it for farmers vs. non-farmers.

Solution. Direct computation:

$$\hat{b} \;=\; \frac{\mathrm{Cov}[(,Y]_i, C_i)}{\mathrm{Var}[(]\,Y_i)} \;=\; \frac{\mathrm{Cov}[(,Y]_i^P + Y_i^T,\, \phi Y_i^P)}{\mathrm{Var}[(]\,Y_i^P) + \mathrm{Var}[(]\,Y_i^T)} \;=\; \phi \cdot \frac{\mathrm{Var}[(]\,Y_i^P)}{\mathrm{Var}[(]\,Y_i^P) + \mathrm{Var}[(]\,Y_i^T)},$$

using $\mathrm{Cov}[(,Y]^P, Y^T) = 0$. Up to the constant factor ϕ (close to 1 in the infinite-horizon limit), the regression slope equals *the share of permanent variance in total income variance.*

Application.

- **Farmers.** Their incomes are dominated by transitory shocks (weather, crop prices), so $\mathrm{Var}[(]\,Y^T)$ is large relative to $\mathrm{Var}[(]\,Y^P)$. The signal-to-total ratio is small, hence $\hat{b}$ is small. The estimated consumption function looks flat, and the "apparent MPC" is low.
- **Non-farmers.** Wage income is much more stable, with $\mathrm{Var}[(]\,Y^T)$ small relative to $\mathrm{Var}[(]\,Y^P)$. The signal-to-total ratio is close to 1, so $\hat{b} \approx \phi \approx 1$. The estimated consumption function looks steep, with apparent MPC close to 1.

Punchline.

- A naive observer comparing the two slopes might conclude that farmers "don't respond" to income while non-farmers respond strongly. The PIH explanation is exactly the opposite: *both* groups respond identically to permanent income changes; the regression slope is small for farmers because their measured income is mostly noise that they correctly ignore.
- The key formula $\hat{b} = \phi \cdot \mathrm{Var}[(]\,Y^P)/[\mathrm{Var}[(]\,Y^P) + \mathrm{Var}[(]\,Y^T)]$ is exam-standard for any cross-section question about apparent MPC vs. structural MPC.

12.1.5 CRRA Consumption with Lognormal Shocks

Question. A consumer has CRRA utility and faces stochastic future consumption. The exam asks for four things in sequence:

(a) Write down the Euler equation.

(b) Assume conditional lognormality of C_{t+1} and rewrite the Euler equation in terms of $\mathbb{E}_t[\ln C_{t+1}]$, $\ln C_t$, σ^2, and primitives.

(c) Show that this implies $\ln C$ follows a random walk with drift.

(d) Interpret how the drift depends on σ^2, in light of the precautionary saving motive.

The economic point is to show that under CRRA (which has $u''' > 0$), uncertainty about future consumption raises expected consumption growth—a smoking-gun signature of precautionary saving that distinguishes CRRA from quadratic utility.

Setup.

- Utility $u(C) = C^{1-\theta}/(1-\theta)$, $\theta > 0$, $\theta \neq 1$. (Marginal utility $u'(C) = C^{-\theta}$.)
- Constant interest rate r, but $\beta(1+r) \neq 1$ in general.
- The conditional distribution of $\ln C_{t+1}$ given time-t information is normal with conditional mean $\mu_t = \mathbb{E}_t[\ln C_{t+1}]$ and conditional variance σ^2.

Solution.

Step 1: CRRA Euler equation. From $u'(C_t) = \beta(1+r)\mathbb{E}_t[u'(C_{t+1})]$:

$$C_t^{-\theta} \;=\; \beta(1+r)\,\mathbb{E}_t\big[C_{t+1}^{-\theta}\big].$$

Step 2: Lognormal substitution. If $\ln C_{t+1} \sim \mathcal{N}(\mu_t, \sigma^2)$ conditional on time-t, then $-\theta \ln C_{t+1} \sim \mathcal{N}(-\theta\mu_t, \theta^2\sigma^2)$, and the lognormal MGF identity $\mathbb{E}[e^x] = e^{\mu+V/2}$ (for $x \sim \mathcal{N}(\mu, V)$) gives

$$\mathbb{E}_t\big[C_{t+1}^{-\theta}\big] \;=\; \mathbb{E}_t\big[e^{-\theta \ln C_{t+1}}\big] \;=\; \exp\Big(-\theta\mu_t + \tfrac{\theta^2\sigma^2}{2}\Big).$$

Substitute into the Euler equation and take logs:

$$-\theta \ln C_t \;=\; \ln\beta + \ln(1+r) - \theta\mu_t + \tfrac{\theta^2\sigma^2}{2}.$$

Step 3: Drifted random walk. Solve for μ_t:

$$\boxed{\mathbb{E}_t[\ln C_{t+1}] \;=\; \ln C_t + \underbrace{\frac{\ln\beta + \ln(1+r)}{\theta}}_{\text{intertemporal substitution}} + \underbrace{\frac{\theta\sigma^2}{2}}_{\text{precautionary}}.}$$

Defining the innovation $u_{t+1} \equiv \ln C_{t+1} - \mathbb{E}_t[\ln C_{t+1}]$ (which has $\mathbb{E}_t[u_{t+1}] = 0$), we have

$$\ln C_{t+1} \;=\; \ln C_t + a + u_{t+1}, \qquad a \equiv \tfrac{\ln\beta(1+r)}{\theta} + \tfrac{\theta\sigma^2}{2},$$

i.e., $\ln C_t$ is a random walk with drift a.

Step 4: Interpretation. The drift has two pieces:

- **Intertemporal-substitution channel** $\frac{\ln \beta(1+r)}{\theta}$: positive when $\beta(1+r) > 1$ (the market rewards saving more than the household discounts the future), negative otherwise. With $\beta(1+r) = 1$, this term vanishes.

- **Precautionary channel** $\frac{\theta\sigma^2}{2}$: *always positive when* $\theta > 0$. A higher σ^2 raises expected consumption growth—the household saves more today (depressing C_t), which mechanically lifts $\mathbb{E}_t[\Delta \ln C_{t+1}]$.

Punchline.

- The lognormal trick ($\mathbb{E}[e^x] = e^{\mu+V/2}$) converts the multiplicative Euler equation into an *additive* relationship in log consumption, exposing the random-walk-with-drift structure.

- The precautionary term $\theta\sigma^2/2$ exists precisely because $u''' > 0$ for CRRA. Under quadratic utility ($u''' = 0$), this term would vanish, and consumption growth would not respond to uncertainty—a smoking gun for distinguishing the two utility classes.

- Memorize the lognormal MGF identity. It is the single technical tool that makes this problem solvable in closed form.

12.1.6 Random-Walk Consumption: Martingale Property and MPC Out of AR(1) Shocks

Question. A consumer follows the Hall-style consumption rule

$$C_t \;=\; \frac{r}{1+r}\big[A_t + H_t\big], \qquad H_t \equiv \sum_{s=0}^{\infty} \frac{\mathbb{E}_t[Y_{t+s}]}{(1+r)^s},$$

where A_t is non-state-contingent assets and H_t is human wealth (the present value of expected income). The exam asks two parts:

(a) Show that $\mathbb{E}_t[C_{t+1}] = C_t$ (consumption is a martingale).

(b) Suppose income follows an AR(1) process $Y_t = \phi Y_{t-1} + u_t$ with i.i.d. innovations u_t. If a unit shock $u_t = 1$ arrives, by how much does consumption C_t change?

The economic content: (a) tests your ability to derive the random-walk property using the law of iterated expectations on human wealth; (b) tests whether you can compute the marginal propensity to consume out of a persistent income shock.

Setup.

- The consumption rule above (you may take this as given; it is derived from the PIH theorem).

- Asset evolution: $A_{t+1} = (1+r)(A_t + Y_t - C_t)$.

- Income process for part (b): $Y_t = \phi Y_{t-1} + u_t$ with $|\phi| < 1$ and $\mathbb{E}_{t-1}[u_t] = 0$.

Solution to (a): Consumption is a martingale.

The key step is computing $\mathbb{E}_t[H_{t+1}]$. By the law of iterated expectations, $\mathbb{E}_t[\mathbb{E}_{t+1}[Y_{t+1+s}]] = \mathbb{E}_t[Y_{t+1+s}]$, so

$$\mathbb{E}_t[H_{t+1}] \;=\; \mathbb{E}_t\left[\sum_{s=0}^{\infty} \frac{\mathbb{E}_{t+1}[Y_{t+1+s}]}{(1+r)^s}\right] \;=\; \sum_{s=0}^{\infty} \frac{\mathbb{E}_t[Y_{t+1+s}]}{(1+r)^s}.$$

Reindex with $s' = s+1$ to align with the definition of H_t:

$$\mathbb{E}_t[H_{t+1}] \;=\; (1+r) \sum_{s'=1}^{\infty} \frac{\mathbb{E}_t[Y_{t+s'}]}{(1+r)^{s'}} \;=\; (1+r)\big(H_t - Y_t\big).$$

(The last equality is H_t minus its $s = 0$ term.) This is the **key lemma**.

Combine with the asset law $\mathbb{E}_t[A_{t+1}] = (1+r)(A_t + Y_t - C_t)$:

$$\mathbb{E}_t[A_{t+1} + H_{t+1}] \;=\; (1+r)(A_t + Y_t - C_t) + (1+r)(H_t - Y_t) \;=\; (1+r)(A_t + H_t - C_t).$$

Therefore

$$\mathbb{E}_t[C_{t+1}] \;=\; \frac{r}{1+r}\,(1+r)(A_t + H_t - C_t) \;=\; r(A_t + H_t) - rC_t.$$

Substitute $r(A_t + H_t) = (1+r)C_t$ from the consumption rule:

$$\mathbb{E}_t[C_{t+1}] \;=\; (1+r)C_t - rC_t \;=\; C_t. \qquad \square$$

Solution to (b): MPC out of an AR(1) innovation.

A unit innovation $u_t = 1$ at time t raises $\mathbb{E}_t[Y_{t+s}]$ by ϕ^s for each $s \geq 0$ (start from $\phi^0 = 1$ at $s = 0$, then ϕ^1 at $s = 1$, etc.). The change in human wealth is a geometric series:

$$\Delta H_t \;=\; \sum_{s=0}^{\infty} \frac{\phi^s}{(1+r)^s} \;=\; \frac{1}{1 - \phi/(1+r)} \;=\; \frac{1+r}{1+r-\phi}.$$

The change in consumption follows from $\Delta C_t = \frac{r}{1+r}\Delta H_t$ (assets A_t are unchanged at the moment of the shock):

$$\boxed{\Delta C_t \;=\; \frac{r}{1+r} \cdot \frac{1+r}{1+r-\phi} \;=\; \frac{r}{1+r-\phi}.}$$

Limit cases.

- $\phi \to 0$ (i.i.d. income): MPC $\to r/(1+r) \approx r$ for small r. The agent treats a one-period windfall as adding only its annuity value to lifetime resources.
- $\phi \to 1$ (permanent shock): MPC $\to 1$. A permanent change in income raises consumption nearly one-for-one—the textbook PIH response.
- Intermediate ϕ: the MPC interpolates between the two limits, reflecting how persistent the shock is expected to be.

Punchline.

- **Random-walk derivation.** The trick is to first establish the lemma $\mathbb{E}_t[H_{t+1}] = (1 + r)(H_t - Y_t)$, then combine with the asset law. Don't try to manipulate $\mathbb{E}_t[C_{t+1}]$ directly without this lemma—it gets messy.
- **MPC formula.** MPC $= r/(1 + r - \phi)$ is monotone increasing in ϕ. It interpolates smoothly between $\sim r$ (i.i.d.) and 1 (permanent), reflecting how much of an income shock is expected to persist into the future.

12.2 Computation Drill

The exam will ask you to write down algorithm steps, not actual code. Each algorithm below should be reproduced from memory: the data structures (grids), the update rule (Bellman or pushforward), the convergence criterion, and—when relevant—the theoretical guarantee (contraction mapping for VFI; fixed point for stationary distribution; market clearing for Aiyagari).

12.2.1 Value Function Iteration (Deterministic Neoclassical Model)

Algorithm: Deterministic VFI

Solving $V(k) = \max_{k'}\{u(Af(k) + (1 - \delta)k - k') + \beta V(k')\}$ on a discretized state space.

Step 1: Discretize the state space. Choose a grid $\mathcal{K} = \{k_1, k_2, \ldots, k_N\}$ for capital. Common choice: $N = 500$ points uniformly on $[k_{\min}, k_{\max}]$ where $[k_{\min}, k_{\max}]$ brackets the steady state k^*. Use a log-spaced grid if precision is needed near $k = 0$ (Inada singularity).

Step 2: Initial guess. $V^{(0)}(k) = 0$ for all $k \in \mathcal{K}$. Any bounded initial guess converges to the same fixed point by the Banach Fixed Point Theorem.

Step 3: Bellman update. For each $k \in \mathcal{K}$, compute

$$V^{(i+1)}(k) \;=\; \max_{k' \in \mathcal{K}\,:\,c>0} \Big\{u\big(Af(k) + (1-\delta)k - k'\big) + \beta V^{(i)}(k')\Big\}.$$

The constraint $c > 0$ requires $k' < Af(k) + (1 - \delta)k$; values violating it are excluded from the maximand. Also record the optimizer at each grid point.

Step 4: Convergence check. Stop when $\max_k |V^{(i+1)}(k) - V^{(i)}(k)| < \varepsilon$ (typically $\varepsilon = 10^{-6}$). Otherwise increment i and return to Step 3. Always set a maximum iteration count (e.g., 1000) to guard against runaway loops.

Step 5: Extract policy. After convergence, the optimal policy is $g(k) = \arg\max_{k'}$ from the final iteration. Apply iteratively to track transition paths from any initial k_0.

Why it works. The Bellman operator T defined by $(TV)(k) = \max_{k'}\{u(\cdot) + \beta V(k')\}$ is a contraction mapping with modulus $\beta < 1$ on the space of bounded continuous functions.

By the Banach Fixed Point Theorem there exists a unique fixed point V^*, and $V^{(i)} \to V^*$ uniformly from *any* bounded initial guess. This is what justifies starting at $V^{(0)} = 0$.

Common pitfalls.

- *Log/CRRA singularity at $c = 0$.* Always enforce $k' < Af(k) + (1-\delta)k$ strictly to avoid $\ln(0) = -\infty$ or $C^{-\theta} \to \infty$.
- *Grid endpoint binding.* If the optimal k' is consistently at the maximum grid value, the grid is too narrow; widen it.
- *Convergence diagnostics.* Plot the maximum change $\max_k |V^{(i+1)} - V^{(i)}|$ across iterations; it should fall geometrically at rate β.

12.2.2 Stochastic VFI for the Real Business Cycle Model

Algorithm: Stochastic VFI for RBC

Solving $V(a,k) = \max_{k'}\{u(c) + \beta\,\mathbb{E}[V(a',k') \mid a]\}$ where a follows an AR(1).

Step 1: Discretize the AR(1) process. If $a' = \rho a + \varepsilon$, $\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon^2)$, use the **Tauchen method**: choose M grid points $\{a^1, \ldots, a^M\}$ spanning ± 3 unconditional standard deviations of a. Compute the $M \times M$ transition matrix Π with entries

$$\pi_{ij} = \Phi\left(\frac{a^j + \Delta/2 - \rho a^i}{\sigma_\varepsilon}\right) - \Phi\left(\frac{a^j - \Delta/2 - \rho a^i}{\sigma_\varepsilon}\right),$$

where Δ is the grid spacing and Φ is the standard normal CDF. As $M \to \infty$, the discrete chain converges to the AR(1).

Step 2: Discretize capital. Choose $\mathcal{K} = \{k_1, \ldots, k_N\}$ as in the deterministic case.

Step 3: Bellman update. For each $(a,k) \in \{a^1, \ldots, a^M\} \times \mathcal{K}$:

$$V^{(i+1)}(a,k) = \max_{k' \in \mathcal{K}\,:\,c>0} \left\{ u\big(e^a f(k) + (1-\delta)k - k'\big) + \beta \sum_{a'} \pi(a' \mid a)\, V^{(i)}(a',k') \right\}.$$

Record the optimizer.

Step 4: Convergence check. Stop when $\max_{a,k} |V^{(i+1)} - V^{(i)}| < \varepsilon$.

Step 5: Simulate. Starting from (a_0, k_0), draw a path $\{a_t\}_{t=1}^T$ from the Markov chain, then iterate $k_{t+1} = g(a_t, k_t)$ using the converged policy. Compute realized $\{y_t = e^{a_t} f(k_t),\ c_t,\ i_t = k_{t+1} - (1-\delta)k_t\}$.

Step 6: Compute moments. Report coefficients of variation: $\mathrm{CV}(x) = \sqrt{\mathrm{Var}[(]x)}/\mathbb{E}(x)$, comparing $\mathrm{CV}(c) < \mathrm{CV}(y) < \mathrm{CV}(i)$ to confirm the model reproduces the canonical investment-volatility fact.

Why CV rather than raw variance. The series $\{c, y, i\}$ have very different mean levels, so raw variance comparisons confound scale with volatility. The coefficient of variation

$\mathrm{CV}(x) = \sqrt{\mathrm{Var}[(] x)}/\mathbb{E}(x)$ is unit-free and is the standard object reported in the RBC literature.

12.2.3 Aiyagari Outer Loop (Three Nested Loops)

Algorithm: Aiyagari Stationary Equilibrium

Find r^* such that the household sector's aggregate savings equal the firm's capital demand.

Step 1: Outer loop initial guess. Set r_0 slightly below $1/\beta - 1$ (since precautionary motives push r^* below the representative-agent benchmark).

Step 2: Compute w and K^{firm} from firm FOCs. For Cobb-Douglas $F(K, L) = K^\alpha L^{1-\alpha}$:

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

Here $L = \mathbb{E}[\exp(z)]$ is the constant aggregate labor supply.

Step 3: Solve household problem (middle loop). Given (r_0, w), run VFI on the discretized (z, k) state space to obtain the policy function $g_k(z, k)$. The Bellman is

$$V(z,k) = \max_{k'} \{u(c) + \beta\, \mathbb{E}[V(z',k') \mid z]\}, \quad c = (1+r_0-\delta)k + we^z - k', \quad k' \geq -B.$$

Step 4: Solve stationary distribution (inner loop). Given g_k, find the fixed point Λ^* of the pushforward operator:

- Initialize $\Lambda^{(0)}$ uniform on the grid.
- Iterate $\Lambda^{(j+1)}(z', k') = \sum_{z,k} \Lambda^{(j)}(z,k) \cdot \pi(z' \mid z) \cdot \mathbf{1}\{g_k(z,k) = k'\}$.
- Off-grid lottery: when $g_k(z, k)$ falls between grid points $k_{j'}$ and $k_{j'+1}$, split the mass linearly: send fraction $1-\theta$ to $k_{j'}$ and θ to $k_{j'+1}$, where θ is the relative position. This preserves total mass exactly.
- Stop when $\|\Lambda^{(j+1)} - \Lambda^{(j)}\| < \varepsilon$.

Step 5: Compute excess demand.

$$X(r_0) \;=\; K^{\text{firm}}(r_0) \;-\; \int g_k(z,k)\, d\Lambda^*(z,k) \;=\; K^{\text{firm}} - K^{\text{HH}}.$$

Step 6: Convergence check on outer loop.

- If $|X(r_0)| < \varepsilon_r$, stop. Equilibrium is $r^* = r_0$.
- If $|X(r_0)| \geq \varepsilon_r$, update via bisection (if monotonicity is known) or relaxation $r_0 \leftarrow r_0 - \kappa \cdot X(r_0)$ for small $\kappa > 0$. Return to Step 2.

Economic interpretation. The outer loop traces out the intersection of the household sector's upward-sloping capital supply $K^{\text{HH}}(r)$ (precautionary saving rises with r) and the firm sector's downward-sloping capital demand $K^{\text{firm}}(r) = L(\alpha/r)^{1/(1-\alpha)}$. At the equilibrium r^*, the two intersect.

The atom at the borrowing constraint. A robust outcome of the algorithm is that the stationary distribution Λ^* has an atom at $k = -B$: the mass of households whose policy $g_k(z,k)$ would set $k' < -B$ in the absence of the constraint and so are forced to $k' = -B$ exactly. This atom is the source of wealth inequality at the bottom of the distribution, distinct from the long right tail generated by persistently lucky high-z households.

12.3 Quick Reference Cheat Sheet

Remark (Formulas to Reproduce From Memory).

Result	Formula
Solow BGP per-worker growth	$g_k = g_A/(1-\alpha)$
Solow BGP total-capital growth	$g_K = g_L + g_A/(1-\alpha)$
Neoclassical Euler	$u'(C_t) = \beta\, u'(C_{t+1})\,[F_K(K_{t+1}, L) + 1 - \delta]$
Hall's RWH (quadratic u, $\beta(1+r) = 1$)	$C_t = \mathbb{E}_t[C_{t+1}]$
CRRA + lognormal drift	$\mathbb{E}_t[\ln C_{t+1}] - \ln C_t = \frac{\ln\beta(1+r)}{\theta} + \frac{\theta\sigma^2}{2}$
PIH consumption rule	$C = \phi \cdot \text{PI}, \quad \phi = r/(1+r)$ (infinite horizon)
PIH MPC out of AR(1) shock	$\text{MPC} = r/(1+r-\phi)$
PIH cross-section regression slope	$\hat{b} = \text{Var}(Y^P)/[\text{Var}(Y^P) + \text{Var}(Y^T)]$
RBC (LQ) policy function	$C_t = \rho K_t + \frac{\rho}{1+\rho-\phi} e_t$
Banach contraction guarantee	$V^{(i+1)} = TV^{(i)}$ converges if T is contraction with modulus < 1

Remark (Common Exam Tricks).

- **Log utility cancels income/sub effects.** When labor income is upfront and the agent saves at random return r, C_1 is independent of the distribution of r—the substitution effect (higher $r \Rightarrow$ save more) and income effect (higher lifetime wealth $\Rightarrow$ consume more) exactly cancel. Recognize this whenever you see log utility plus saving.

- **Jensen's inequality for mean-preserving spreads.** If a function f is convex (resp. concave), an MPS in the random variable raises (resp. lowers) $\mathbb{E}[f]$. To apply this to a comparative-static question, identify the relevant f inside the FOC and check the sign of f'' by direct differentiation.

- **Guess-and-verify for linear-quadratic problems.** The closed-form solution to LQ

Bellman problems is always linear in the state. Guess $C_t = \alpha + \beta K_t + \gamma e_t$, substitute into the resource constraint to express K_{t+1} in terms of K_t, e_t, then substitute into the Euler equation and match coefficients. The constant (α), capital-loading (β), and shock-loading (γ) are pinned down independently.

- **Lognormal MGF.** If $x \sim \mathcal{N}(\mu, V)$, then $\mathbb{E}[e^x] = e^{\mu+V/2}$. Use this whenever the Euler equation contains $\mathbb{E}[C^{-\theta}]$ and $\ln C$ is conditionally normal.
- **Random-walk ⇒ AR(1) MPC.** The PIH consumption response to an AR(1) income innovation with persistence ϕ is $r/(1+r-\phi)$. Limits: $\phi = 0 \Rightarrow$ MPC $\approx r$; $\phi = 1 \Rightarrow$ MPC $\to 1$.
- **Computational algorithms have a common skeleton.** Discretize state, guess initial value, update with Bellman or pushforward, check convergence, extract policy / aggregate. Variations: deterministic vs. stochastic (add a discrete-shock dimension and an expectation in the Bellman); Aiyagari adds an outer loop on r.

Part III

Problem Sets and Solutions

Problem Set 1

Problems

Problem 1.1: Risk Sharing in Complete Markets

Consider a one-period economy. There are two countries (or agents), $i = 1, 2$, and two goods, $j = A, B$. There is a state of the world indexed by $s \in S$, occurring with probability $\pi(s) > 0$. In state s, country 1 is endowed with $e^A(s)$ units of good A and country 2 with $e^B(s)$ units of good B (each country has zero endowment of the other good). Each country's representative agent has utility

$$u\big(c_i^A(s),\, c_i^B(s)\big) \;=\; \frac{(c_i^A(s)\, c_i^B(s))^{(1-\rho)/2}}{1-\rho}.$$

1. Solve for the Pareto optimal allocations $\{c_1^A, c_1^B, c_2^A, c_2^B\}$. Take the Pareto weight for country 1 as λ and for country 2 as $1-\lambda$.
2. Suppose there are no financial markets, but countries can trade goods after the state is realized. Let $p(s)$ denote the relative price of good B in terms of good A.
   (a) Define the competitive equilibrium.
   (b) Solve for the competitive equilibrium.
   (c) Show that the competitive equilibrium allocation is Pareto efficient.
   (d) What happens to the relative price of good B in terms of A in states where $e^A(s)/e^B(s)$ is low? Explain.
3. Explain what allows countries to deal with risk sharing in the competitive equilibrium with no financial markets.

Problem 1.2: Dynamic Programming

A representative consumer lives forever. The final good can be consumed or saved as capital. The consumer is endowed with one unit of labor each period; some is consumed as leisure x_t and some supplied as labor ℓ_t. Initial capital k_0 is given; capital depreciates at rate δ. Preferences are $\sum_{t=0}^{\infty} \beta^t u(c_t, x_t)$. A representative firm produces a single final good using technology $f(k_t, \ell_t)$.

1. Define a sequential market equilibrium.

2. Obtain the first-order conditions and explain how to find the equilibrium allocations.

3. Write down the consumer's dynamic programming problem (Bellman equation).

4. State the parameter assumptions guaranteeing a solution to the Bellman equation. Explain why each is needed.

5. Obtain the FOCs of the dynamic programming problem and explain how to find the allocations.

6. Are the allocations from (2) and (5) equivalent? Show your work.

7. Write down the dynamic programming problem for the social planner.

8. Find the FOCs of the social planner's problem, and show that the planner's allocations coincide with the competitive allocations from (2).

Problem 1.3: Heterogeneous Discounting

Each period $t \geq 0$ a stochastic event $s_t \in S$ is realized. The unconditional probability of a history $s^t = (s_0, s_1, \dots, s_t)$ is $\pi(s^t)$. There are I consumers; consumer i is endowed with $y^i(s^t)$ of the single good in history s^t. Each consumer has utility

$$U(c^i) \;=\; \sum_{t=0}^{\infty} \sum_{s^t} \beta_i^t \, \pi(s^t) \, u\big(c^i(s^t)\big),$$

with consumer-specific discount factors $1 > \beta_1 \geq \beta_2 \geq \cdots \geq \beta_I > 0$.

1. Write down the Pareto problem.

2. Obtain the FOCs for the Pareto problem.

3. How do they differ from the FOCs of the problem with a common discount factor?

4. Show that consumption of i, where $\beta_i < \beta_1$, is decreasing over time.

5. Show that the most patient consumer eventually consumes everything.

Solutions

Problem 1.1: Risk Sharing in Complete Markets

(1) Pareto optimal allocations. The planner solves

$$\max_{\{c_1^A, c_1^B, c_2^A, c_2^B\}_s} \quad \sum_s \pi(s) \left[\lambda \frac{(c_1^A c_1^B)^\gamma}{1-\rho} + (1-\lambda) \frac{(c_2^A c_2^B)^\gamma}{1-\rho} \right]$$
$$\text{s.t.} \quad c_1^A(s) + c_2^A(s) = e^A(s), \qquad c_1^B(s) + c_2^B(s) = e^B(s),$$

where $\gamma \equiv (1-\rho)/2$. Letting $\mu_A(s), \mu_B(s)$ denote the multipliers, the FOCs imply

$$\frac{c_i^B(s)}{c_i^A(s)} = \frac{\mu_A(s)}{\mu_B(s)} \quad (i = 1, 2) \implies \frac{c_1^B(s)}{c_1^A(s)} = \frac{c_2^B(s)}{c_2^A(s)},$$

so both agents consume the two goods in the same proportion in every state. Equating the shadow prices $\mu_A(s)$ from agents 1 and 2 then yields

$$\frac{c_1^A(s)}{c_2^A(s)} = \left(\frac{\lambda}{1-\lambda} \right)^{1/\rho} \equiv \lambda_0.$$

Combining with the resource constraint:

$$\boxed{c_1^A(s) = \frac{\lambda_0}{1+\lambda_0} e^A(s), \quad c_2^A(s) = \frac{1}{1+\lambda_0} e^A(s), \quad \text{and analogously for } B.}$$

Two key features: (i) consumption shares depend on the Pareto weight λ but not on the state, so each agent consumes a state-invariant share of the aggregate; (ii) the allocation depends on aggregate endowments $\{e^A(s), e^B(s)\}$ but not on which country owns each good.

(2a) Competitive equilibrium definition. A competitive equilibrium is a price sequence $\{p(s)\}_{s \in S}$ and an allocation $\{c_i^A(s), c_i^B(s)\}_{i,s}$ such that:

- *Country 1's problem:* max $\pi(s)(c_1^A c_1^B)^\gamma/(1-\rho)$ subject to $c_1^A(s) + p(s)c_1^B(s) \leq e^A(s)$.
- *Country 2's problem:* max $\pi(s)(c_2^A c_2^B)^\gamma/(1-\rho)$ subject to $c_2^A(s) + p(s)c_2^B(s) \leq p(s)e^B(s)$.
- *Market clearing:* $c_1^A(s) + c_2^A(s) = e^A(s)$ and $c_1^B(s) + c_2^B(s) = e^B(s)$ for every s.

(2b) Solving for the equilibrium. Since utility is a monotonic transformation of Cobb–Douglas with equal exponents, each agent spends half their wealth on each good. For country 1 (wealth $e^A(s)$), this gives $c_1^A(s) = \frac{1}{2}e^A(s)$ and $c_1^B(s) = e^A(s)/[2p(s)]$. For country 2 (wealth $p(s)e^B(s)$), $c_2^A(s) = \frac{1}{2}p(s)e^B(s)$ and $c_2^B(s) = \frac{1}{2}e^B(s)$. Imposing market clearing for good A:

$$\tfrac{1}{2}e^A(s) + \tfrac{1}{2}p(s)e^B(s) = e^A(s) \implies \boxed{p(s) = \frac{e^A(s)}{e^B(s)}.}$$

Substituting back, $c_i^A(s) = \frac{1}{2}e^A(s)$ and $c_i^B(s) = \frac{1}{2}e^B(s)$ for both i.

(2c) Pareto efficiency. The competitive equilibrium delivers $c_1^A/c_2^A = c_1^B/c_2^B = 1$, which is exactly the Pareto allocation with $\lambda_0 = 1$, i.e., $\lambda = 1/2$. So the competitive allocation is Pareto efficient (with equal weights) by the First Welfare Theorem.

(2d) Price response to skewed endowments. If $e^A(s)/e^B(s)$ is low, good A is scarce relative to good B, so the relative price of B in terms of A falls: $p(s) = e^A(s)/e^B(s)$ is small. Intuitively, country 1 (the holder of the abundant good A) has weak demand for B, while country 2 (holder of the scarce B) is willing to part with B cheaply to acquire the more valuable A.

(3) The risk-sharing mechanism. Even without explicit financial markets, the spot relative price $p(s)$ moves with the state and acts as a state-contingent claim. When a country's endowment is high in some state, the corresponding good becomes cheap, eroding the country's terms of trade and limiting how much extra consumption it can buy. Conversely, when the endowment is low, the relative price of its good rises, propping up its purchasing power. The price mechanism therefore implements the same risk-sharing outcome that a complete state-contingent financial market would, as long as goods can be traded after the state realizes.

Problem 1.2: Dynamic Programming

(1) Sequential market equilibrium. Households take prices $\{w_t, r_t\}$ as given and solve

$$\max_{\{c_t, x_t, \ell_t, k_{t+1}\}} \sum_{t=0}^{\infty} \beta^t u(c_t, x_t) \text{ s.t. } c_t + k_{t+1} = w_t \ell_t + (r_t + 1 - \delta) k_t, \;\; x_t + \ell_t = 1.$$

Firms take prices and solve $\max_{k_t, \ell_t} f(k_t, \ell_t) - r_t k_t - w_t \ell_t$ each period. An equilibrium is a sequence of prices and allocations such that both problems are solved and markets clear: $c_t + k_{t+1} - (1-\delta)k_t = f(k_t, \ell_t)$, capital and labor markets clear period by period.

(2) FOCs and equilibrium allocations. Household FOCs (with u_c, u_x partials):

$$\underbrace{\frac{u_x(c_t, x_t)}{u_c(c_t, x_t)} \;=\; w_t,}_{\text{intratemporal MRS}} \qquad \underbrace{u_c(c_t, x_t) \;=\; \beta\, u_c(c_{t+1}, x_{t+1})(r_{t+1} + 1 - \delta)}_{\text{Euler equation}}.$$

Firm FOCs: $r_t = f_k(k_t, \ell_t)$, $w_t = f_\ell(k_t, \ell_t)$. Substituting prices into the household conditions and combining with market clearing yields a sequence of equations that pin down $\{c_t, x_t, k_{t+1}\}$ recursively, given k_0.

(3) The Bellman equation.

$$V(k) \;=\; \max_{c, x, \ell, k'} \{\, u(c, x) + \beta V(k') \,\} \quad \text{s.t.} \quad c + k' = w(k)\ell + (r(k) + 1 - \delta)k, \;\; x + \ell = 1.$$

(4) Assumptions for a solution.

- *Bounded state space* (or compact-valued constraint correspondence): ensures the maximum exists each period.
- *Continuity and concavity* of u and f: required for the contraction-mapping argument and uniqueness of the maximizer.
- *Discounting* $\beta \in (0, 1)$: makes the Bellman operator a contraction with modulus β, ensuring a unique fixed point by the Banach Fixed-Point Theorem.

- *Standard monotonicity* ($u_c, f_k > 0$) and concavity ($u_{cc}, f_{kk} < 0$): ensure interior solutions and a well-defined value function.
- *Inada conditions* (optional): guarantee interior $c, k > 0$ solutions and rule out corners.

(5) FOCs of the dynamic program. Differentiating the RHS of the Bellman:

$$u_c(c,x) \;=\; \beta\, V'(k') \cdot 1, \qquad u_x(c,x) \;=\; u_c(c,x)\, w(k).$$

Combined with the envelope condition $V'(k) = u_c(c,x)(r(k)+1-\delta)$, this recovers the Euler equation $u_c(c_t, x_t) = \beta\, u_c(c_{t+1}, x_{t+1})(r_{t+1} + 1 - \delta)$ and the intratemporal MRS condition.

(6) Equivalence of (2) and (5). The intratemporal MRS condition and the Euler equation derived recursively in (5) are identical to those obtained directly from the sequence problem in (2). Both give the same system, so the allocations coincide. The Bellman approach simply recasts the infinite-dimensional problem $\{c_t, x_t, k_{t+1}\}_{t=0}^{\infty}$ as a sequence of two-period decisions linked by the value function—an operationally simpler but mathematically equivalent formulation.

(7) Social planner's dynamic program. The planner maximizes household utility subject to the resource constraint:

$$V^{SP}(k) \;=\; \max_{c,x,\ell,k'} \left\{ u(c,x) + \beta V^{SP}(k') \right\} \quad \text{s.t.} \quad c + k' \;=\; f(k,\ell) + (1-\delta)k, \;\; x + \ell = 1.$$

No prices appear: the planner allocates output directly.

(8) Planner FOCs and equivalence. The FOCs are

$$u_c \;=\; \beta\, u_c'\,(f_k' + 1 - \delta), \qquad u_x \;=\; u_c\, f_\ell.$$

These are exactly the competitive Euler equation and intratemporal condition once we impose $r_t = f_k$ and $w_t = f_\ell$ from the firm's FOCs. Hence the planner's allocations coincide with the competitive ones, illustrating the First Welfare Theorem in this setting.

Problem 1.3: Heterogeneous Discounting

(1) Pareto problem. For weights $\{\eta_i > 0\}_{i=1}^{I}$ summing to one, the planner solves

$$\begin{aligned} \max_{\{c^i(s^t)\}} \quad & \sum_{i=1}^{I} \eta_i\, U(c^i) \;=\; \sum_i \eta_i \sum_{t,s^t} \beta_i^t\, \pi(s^t)\, u(c^i(s^t)) \\ \text{s.t.} \quad & \sum_i c^i(s^t) \;\le\; \sum_i y^i(s^t) \quad \forall (t, s^t). \end{aligned}$$

(2) Pareto FOCs. Let $\theta(s^t)$ multiply the resource constraint. The FOC for $c^i(s^t)$ is

$$\eta_i\, \beta_i^t\, \pi(s^t)\, u'(c^i(s^t)) \;=\; \theta(s^t),$$

which yields the cross-agent ratio

$$\boxed{\frac{u'(c^i(s^t))}{u'(c^j(s^t))} = \frac{\eta_j}{\eta_i} \cdot \left(\frac{\beta_j}{\beta_i}\right)^t.}$$

(3) Difference from common-discount case. With a common β, the ratio $u'(c^i)/u'(c^j) = \eta_j/\eta_i$ is constant over time and across states; perfect risk sharing means each agent consumes a constant share of the aggregate. With heterogeneous β_i, the ratio carries an additional $(\beta_j/\beta_i)^t$ factor that grows or shrinks geometrically over time, breaking the time-invariance of consumption shares.

(4) Consumption of less-patient agents falls over time. Take $i \neq 1$, so $\beta_i < \beta_1$. The ratio

$$\frac{u'(c_t^i)}{u'(c_t^1)} = \frac{\eta_1}{\eta_i} \cdot \left(\frac{\beta_1}{\beta_i}\right)^t$$

grows without bound as $t \to \infty$ since $\beta_1/\beta_i > 1$. Because u' is strictly decreasing, this means $c_t^i/c_t^1 \to 0$. To see that c_t^i itself is decreasing for sufficiently large t (and not merely as a share), use the resource constraint and the fact that aggregate endowment is bounded: when c_t^1 rises and total $\sum_i c_t^i$ remains bounded, c_t^i must eventually fall.

(5) The most patient consumer dominates. Letting $t \to \infty$ in the ratio above, $c_t^i/c_t^1 \to 0$ for every $i \neq 1$. Combined with the resource constraint $\sum_i c_t^i = Y_t$ (bounded), this forces $c_t^i \to 0$ for $i \neq 1$ and hence $c_t^1 \to Y_t$. The most patient consumer (smallest effective discount, β_1) eventually consumes the entire aggregate endowment.

Remark (Why patience dominates).

> Heterogeneous β generates a force absent in the standard model: the planner's marginal cost of giving consumption to agent i at date t is $\eta_i\beta_i^t$, which decays faster for impatient agents. The planner therefore optimally postpones their consumption indefinitely. In equilibrium without the planner—when agents trade in complete markets—this same logic implies that the most patient agent buys up all assets in the long run. The heterogeneous-discount economy is therefore an extreme case of the inequality dynamics studied in modern HANK models: even with identical income processes and complete insurance, ex-ante differences in patience generate unbounded long-run wealth divergence.

Problem Set 2

Problems

Problem 2.1: CARA Utility in Incomplete Markets

Consider the partial-equilibrium CARA model with incomplete markets. The agent has access to a one-period risk-free bond paying $R = 1 + r$, and CARA utility $u(c) = -\frac{1}{\gamma}e^{-\gamma c}$. The endowment follows an AR(1) process $y(s') = \phi y(s) + (1-\phi)\bar{y} + \varepsilon(s')$, with $\varepsilon \sim \mathcal{N}(0, \sigma^2)$ and $|\phi| < 1$. The agent is unconstrained in borrowing.

(a) Write down the recursive problem in cash-in-hand $x = y(s) + Ra(s)$, and obtain the FOC and envelope condition.

(b) Use your knowledge of the consumption function and cash-in-hand evolution to derive an expression for $c(x', s') - c(x, s)$, where $x' = R\,a(x, s) + y(s')$.

(c) Show that consumption is a random walk with drift. How should consumption respond to predictable changes in future income?

(d) Show that in general equilibrium, individual consumption is a random walk *without* drift.

Now consider the same CARA environment but where individual endowment splits into an aggregate component and an idiosyncratic component:

$$y^i(s') \;=\; w^i(s) + y(s),$$

with $y(s') = \phi_1 y(s) + (1 - \phi_1)\bar{y} + \varepsilon(s')$ and $w^i(s') = \phi_2 w^i(s) + \eta^i(s')$, both shocks normal i.i.d. with variances σ_ε^2 and σ_η^2.

(e) Guess that the log of the value function is linear in cash-in-hand x, the aggregate shock $y(s)$, and the idiosyncratic shock $w^i(s)$. Solve for the consumption function and the evolution of cash-in-hand.

(f) Derive a formula for $c(x', s') - c(x, s)$. Discuss how consumption responds to the aggregate vs. idiosyncratic shock, and compare to part (b).

(g) Solve for the general equilibrium condition that the interest rate must satisfy. Explain the difference with the equation in class.

(h) In general equilibrium, derive the equation for the evolution of cash-in-hand. How does this differ from the case with idiosyncratic shocks only?

Problem 2.2: Comparing Complete and Incomplete Market Savings

Assume the simple two-period model from class.

1. Write down the incomplete-market problem and derive the FOCs.
2. Write down the complete-market problem and derive the FOCs.
3. Adjust the complete-market problem to make it directly comparable to the incomplete market. Write down the new problem and FOC.
4. Assuming $u''' > 0$, show that assets under complete markets are greater than assets under incomplete markets.

Problem 2.3: Small Open Economy with Incomplete Markets

Consider a small open economy that does not affect the world interest rate. The endowment satisfies $y_{t+1} = y_t(1 + g_t)$, where $g_t = g$ with probability $\pi > 1/2$ and $g_t = -g$ with probability $1 - \pi$, and $g = 2\pi - 1$. Shocks are i.i.d. across t and independent of everything else.

The representative agent saves in a riskless bond paying R abroad, with $a_t \geq \bar{a} < 0$ and $a_0 = 0$. Utility is $u(c) = \ln c$, and $\beta = 1/R$.

1. Write the Bellman equation.
2. Obtain the FOCs and the Euler equation.
3. Show that $c_t = y_t$ is the optimal consumption plan. What is the optimal asset path?
4. How does this differ from the complete-markets allocation? Describe the expected complete-market allocation in words.
5. Are there any precautionary savings? Explain.

Solutions

Problem 2.1: CARA Utility in Incomplete Markets

(a) Recursive problem and FOC. Define cash-in-hand $x = y(s) + Ra(s)$. The Bellman equation is

$$V(x,s) \;=\; \max_c \Big\{ -\tfrac{1}{\gamma}e^{-\gamma c} \;+\; \beta\,\mathbb{E}\big[V(x',s') \mid s\big] \Big\} \quad \text{s.t.} \quad x' = R(x-c) + y(s').$$

The FOC and envelope conditions are

$$e^{-\gamma c} \;=\; \beta R\,\mathbb{E}[V_x(x',s') \mid s], \qquad V_x(x,s) \;=\; e^{-\gamma c}.$$

Combining yields the Euler equation $1 = \beta R\,\mathbb{E}[e^{-\gamma(c'-c)} \mid s]$.

(b) Consumption change. Guess $V(x,y) = -A\exp(-Bx - Dy)$ and solve via matching. Setting $\Psi \equiv (R-1)/(R-\phi)$, the consumption function is

$$c(x,y) \;=\; \frac{R-1}{R}x + \frac{R-1}{R-\phi}\,\frac{\phi}{R}\,y + \text{const} \;=\; (R-1)a + \Psi y + \text{const}.$$

Using the budget constraint $a' = x - c$ and the AR(1) for y, after simplification:

$$\boxed{c(x',y') - c(x,y) \;=\; \mu_0 + \Psi\,\varepsilon', \quad \mu_0 \;=\; \tfrac{1}{\gamma}\ln(\beta R) + \tfrac{\gamma\sigma^2}{2}\,\Psi^2.}$$

(c) Random walk with drift. The expression above shows

$$\mathbb{E}_t[c_{t+1}] \;=\; c_t + \mu_0,$$

i.e., consumption is a random walk with drift μ_0. The change Δc_{t+1} depends only on the constant drift and the unforecastable shock ε', so consumption does *not* respond to predictable changes in income—a strong testable implication of the rational-expectations household.

(d) General equilibrium without drift. In general equilibrium, bond market clearing requires $\int a_{i,t}\,di = 0$. Goods market clearing combined with constant aggregate endowment forces $\Delta C_t = \int (c_{i,t+1} - c_{i,t})\,di = 0$. From part (b), averaging across agents (idiosyncratic ε averages to zero by LLN) gives

$$0 \;=\; \mu_0(R) \;=\; \tfrac{1}{\gamma}\ln(\beta R) + \tfrac{\gamma\sigma^2}{2}\,\Psi^2(R).$$

This pins down the equilibrium interest rate

$$R^* \;=\; \frac{1}{\beta}\exp\Big(-\tfrac{\gamma^2\sigma^2}{2}\,\Psi^2\Big) \;<\; \frac{1}{\beta},$$

and individual consumption obeys $c_{t+1} = c_t + \Psi\,\varepsilon_{t+1}$—a pure random walk.

(e) Two-shock value function. Guess $V(x,y,w) = -A\exp(-Bx - Dy - Fw)$. Matching coefficients (analogously to part (b)) yields

$$B \;=\; \gamma\,\frac{R-1}{R}, \quad D \;=\; B\,\frac{\phi_1}{R-\phi_1}, \quad F \;=\; B\,\frac{\phi_2}{R-\phi_2}.$$

The consumption function is

$$c(a,y,w) \;=\; (R-1)a + \frac{R-1}{R-\phi_1}\,y + \frac{R-1}{R-\phi_2}\,w + \text{const.}$$

Cash-in-hand evolves via $x' = y' + w' + R(x-c)$.

(f) Consumption response to two shocks. Define $\Psi_1 = (R-1)/(R-\phi_1)$ and $\Psi_2 = (R-1)/(R-\phi_2)$. After substitution,

$$\Delta c_{t+1} \;=\; \mu_0 + \Psi_1\,\varepsilon_{t+1} + \Psi_2\,\eta^i_{t+1}.$$

Both shocks pass through to consumption with their respective Ψ coefficients. Since aggregate y has higher persistence ($\phi_1 \approx 1$ in calibration), $\Psi_1 \approx 1$ (full pass-through), while idiosyncratic w^i may have lower persistence, giving smaller Ψ_2.

(g) GE interest rate condition. The GE condition now requires the drift to vanish accounting for both shock variances:

$$\boxed{\tfrac{1}{\gamma}\ln(\beta R) + \tfrac{\gamma\sigma_\varepsilon^2}{2}\,\Psi_1^2 + \tfrac{\gamma\sigma_\eta^2}{2}\,\Psi_2^2 \;=\; 0.}$$

Compared to the single-shock case, the equilibrium R^* is lower: the additional aggregate uncertainty (variance σ_ε^2) generates extra precautionary saving, depressing the interest rate further.

(h) Cash-in-hand evolution in GE. Aggregating individual cash-in-hand: $\int x_{i,t}\,di = \int y_t\,di + \int w_t^i\,di + R\int a_{i,t}\,di$. By LLN $\int w_t^i\,di = 0$, and bond market clearing gives $\int a_{i,t}\,di = 0$, so $\int x_{i,t}\,di = y_t$ in equilibrium. Compared to the idiosyncratic-only case (constant aggregate), here y_t fluctuates with the aggregate AR(1) process, and so does aggregate cash-in-hand.

Problem 2.2: Comparing Complete and Incomplete Market Savings

(1) Incomplete markets. The agent solves

$$\max_a\; u(y_0 - a) + \beta\sum_s \pi(s)\,u(y(s) + Ra),$$

yielding the FOC

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

(2) Complete markets. With state-contingent assets $a(s)$,

$$\max_{c_0,\{a(s)\}} u(c_0) + \beta\sum_s \pi(s)\,u(y(s) + a(s)) \;\text{ s.t. }\; c_0 + \sum_s Q(s)\,a(s) = y_0.$$

FOCs: $u'(c_0)Q(s) = \beta\pi(s)\,u'(y(s) + a(s))$, equivalent to $Q(s) = \beta\pi(s)\,u'(c(s))/u'(c_0)$.

(3) Adjusted complete-market problem. Under $\beta = 1/R$ and assuming complete markets deliver perfect smoothing (so $c(s) = c_0$ across states), state prices simplify to $Q(s) = \beta\pi(s)$. Define $\hat{a} = \sum_s Q(s)a(s) = \beta \sum_s \pi(s)a(s)$. Aggregating the complete-market budget constraint and using $c(s) = c_0$:

$$\max_{\hat{a}} \ u(y_0 - \hat{a}) + \beta\, u\big(\mathbb{E}[y(s)] + R\hat{a}\big).$$

The FOC becomes

$$u'(y_0 - \hat{a}) \ = \ \beta R\, u'\big(\mathbb{E}[y(s)] + R\hat{a}\big).$$

(4) Comparison. The incomplete-market FOC has $\beta R\,\mathbb{E}[u'(y(s) + Ra)]$ on the RHS; the complete-market FOC has $\beta R\, u'(\mathbb{E}[y(s)] + R\hat{a})$. With $u''' > 0$ (prudence), u' is convex, so by Jensen,

$$\mathbb{E}[u'(y(s) + Ra)] \ > \ u'(\mathbb{E}[y(s)] + Ra),$$

i.e., expected marginal utility tomorrow is higher under uncertainty. To restore the FOC equality at the incomplete-market optimum, we need $u'(y_0 - a)$ to be larger, so a must be *larger* than $\hat{a}$ would yield. Hence

$$\boxed{a \ > \ \hat{a},}$$

i.e., agents save more under incomplete markets than under complete markets. This excess saving is the **precautionary saving** motive: facing uninsurable risk, agents accumulate a buffer.

Problem 2.3: Small Open Economy

(1) Bellman equation. With state (y, a) (cash-in-hand reduces to $y(1+g)+Ra$; the only Markov-relevant state is y):

$$V(y, a) \ = \ \max_{c,a'} \big\{ \ln c + \beta\,\mathbb{E}\big[V(y', a') \mid y\big] \big\} \ \text{s.t.} \ \ c + a' = y + Ra, \ \ a' \geq \bar{a}.$$

With $y' = y(1 + g')$, where $g' = \pm g$ with probabilities $\pi, 1 - \pi$.

(2) FOCs and Euler equation.

$$\frac{1}{c} \ = \ \beta R\,\mathbb{E}\left[\frac{1}{c'} \,\middle|\, y\right] \ + \ \mu,$$

where $\mu \geq 0$ is the multiplier on $a' \geq \bar{a}$. With $\beta R = 1$ and assuming the constraint slack:

$$\boxed{\frac{1}{c_t} \ = \ \mathbb{E}_t\left[\frac{1}{c_{t+1}}\right].}$$

(3) Verifying $c_t = y_t$. Conjecture $c_t = y_t$ (and hence $a_t = 0$ for all t, consistent with $a_0 = 0$). Then $c_{t+1} = y_{t+1} = y_t(1 + g')$, and the Euler equation becomes

$$\frac{1}{y_t} \ \overset{?}{=} \ \mathbb{E}_t\left[\frac{1}{y_t(1 + g')}\right] \ = \ \frac{1}{y_t}\,\mathbb{E}_t\left[\frac{1}{1 + g'}\right].$$

The required condition is $\mathbb{E}[1/(1+g')] = 1$. Compute:

$$\mathbb{E}\left[\frac{1}{1+g'}\right] = \frac{\pi}{1+g} + \frac{1-\pi}{1-g} = \frac{\pi(1-g) + (1-\pi)(1+g)}{1-g^2} = \frac{1-g(2\pi-1)}{1-g^2}.$$

Using $g = 2\pi - 1$, the numerator becomes $1 - g \cdot g = 1 - g^2$, equal to the denominator. So $\mathbb{E}[1/(1+g')] = 1$ exactly, verifying the conjecture. **Optimal allocation**: $c_t = y_t$, $a_t = 0$ for all t.

(4) Comparison to complete markets. Under complete markets, the agent purchases state-contingent claims to perfectly smooth consumption: c_t would be constant and equal to permanent income $\mathbb{E}_0[y_t]$ in expectation. In particular, the consumer would borrow in low states and lend in high states. Under the incomplete-market regime here, no such smoothing is possible—consumption tracks endowment one-for-one, and the agent cannot shift resources across states.

(5) No precautionary savings. Despite uninsurable income risk, the agent holds zero assets at every date. The reason: with $\beta R = 1$ and log utility, the calibration of $g = 2\pi - 1$ is precisely the condition under which the certainty-equivalent income equals the actual income realization in expected log terms. Log utility delivers $u''' > 0$ but the specific income process turns the precautionary force "off." Moreover, since $a_t = 0 > \bar{a}$, the borrowing constraint never binds, so there is no constraint-induced saving either.

Remark (The Calibration Trick).

> The condition $g = 2\pi - 1$ implies $\mathbb{E}[(1+g')^{-1}] = 1$, which is exactly the multiplicative analog of $\mathbb{E}[g'] = 0$ that would arise in an additive AR(1) model. This is a deliberate calibration that makes the model tractable: it kills any precautionary saving motive and lets the agent stay at $a = 0$ forever. With a different calibration (e.g., g chosen independently of π), the agent would generally accumulate or decumulate buffer stock, and the closed-form solution would break.

Problem Set 3

Problems

Problem 3.1: One-Sided Lack of Commitment with Markov Process

A risk-neutral money lender (ML) borrows and lends abroad at $R = \beta^{-1}$ and is fully committed. The agent's preferences are

$$\mathbb{E}_{-1} \sum_{t=0}^{\infty} \beta^t u(c(s^t)),$$

where u is increasing, strictly concave, with $u'(0) = +\infty$. The endowment follows a discrete Markov chain with $\Pi_{s,s'} = \Pr(()\, y_{t+1} = y_{s'} \mid y_t = y_s)$, and $y_{s+1} > y_s$. The ML offers a contract minimizing expected discounted cost $C(\cdot)$ of delivering value v to the household.

1. What are the state variables of the ML's recursive problem? Explain why each is needed.
2. Assume both parties fully commit. Write the recursive problem of the lender, where $C(\cdot)$ is the minimum expected discounted cost.
3. Write the Lagrangian. Obtain the FOC and envelope conditions. What can you say about c and the continuation value w?
4. Assume $C(v, s) = \phi(v) + \psi(s)$ separable. What does this imply for the optimal contract?
5. Let $v^a(s)$ be the autarky value conditional on y_s. Obtain its Bellman equation.

Now assume the consumer has limited liability: if they default, they consume their endowment forever.

6. Write the recursive problem of the ML under limited liability.
7. Define the Lagrangian; obtain FOC and envelope conditions.
8. Describe the optimal contract as a function of the Lagrange multiplier on the participation constraint.
9. In words, describe the intuition behind the optimal contract.
10. What is the cost of offering the contract (4) to the ML? Is it self-enforcing?

Problem 3.2: Bulow–Rogoff Contracts

A country with i.i.d. endowment $y(s) \in [y_1, \ldots, y_N]$ trades with a risk-neutral foreigner with discount factor $1/R$ $(R > 1)$. The country has utility $u(c)$, discount β.

1. Assume full commitment for both. Write down the problem defining the first-best efficient contracts.

2. Obtain the FOC and write the Euler equation.

3. What value of β makes the first-best contract stationary? Can we index contracts by $\bar{c}$?

4. There exist N Arrow securities, asset i paying 1 unit if state i is realized. What is the price of asset i?

5. Suppose the country can default. After default, it cannot borrow but can buy positive positions in the assets in (4). The first-best $\bar{c}$ is in place.

   (a) Show that if $\bar{c} < (1-\beta)y_N + \beta\mathbb{E}[y]$, then after the highest endowment y_N, the country is a debtor.

   (b) Does this imply the country is a debtor at $t = 0$ before the first realization?

   (c) Show explicitly (without invoking Bulow–Rogoff) that under the same condition, after the first y_N shock, the country is strictly better off defaulting and saving in the Arrow securities.

Solutions

Problem 3.1: One-Sided Lack of Commitment with Markov Process

(1) State variables. The ML's recursive problem requires two state variables:

- *Promised utility* v: summarizes all future obligations to the consumer. Without it, the ML cannot keep track of past promises in a recursive formulation.
- *Current endowment realization* y_s: with a Markov endowment process, the conditional distribution of future endowments depends on the current state. Hence s is needed to forecast future income and contracts.

(2) Recursive problem (full commitment).

$$C(v,s) \;=\; \min_{c,w(s')} \{\, c - y_s + \beta\,\mathbb{E}[C(w(s'),s') \mid s]\,\}$$

subject to the promise-keeping constraint

$$u(c) + \beta\,\mathbb{E}[w(s') \mid s] \;\geq\; v.$$

(3) Lagrangian and FOCs. Let μ be the multiplier on promise-keeping. The Lagrangian is

$$\mathcal{L} \;=\; c - y_s + \beta\mathbb{E}[C(w,s') \mid s] + \mu\,[v - u(c) - \beta\,\mathbb{E}[w \mid s]].$$

FOCs:

$$1 \;=\; \mu\,u'(c), \qquad \beta\pi(s' \mid s)\,C_w(w(s'),s') \;=\; \mu\beta\pi(s' \mid s).$$

Envelope: $C_v(v,s) = \mu$. Combining, $C_w(w(s'),s') = \mu = C_v(v,s)$, so $w(s') = v$ for every s' (using strict concavity of C in v). Likewise c is constant across states. The consumer's continuation value is constant; consumption is fully smoothed.

(4) Separable cost. If $C(v,s) = \phi(v) + \psi(s)$, then $C_v(v,s) = \phi'(v)$ depends only on v, not s. The optimality condition $C_v(v,s) = C_w(w(s'),s')$ implies $\phi'(v) = \phi'(w(s'))$, hence $w(s') = v$ for all s'. The optimal contract delivers a state-invariant continuation value, and consumption c depends only on v (through $\phi'(v) = 1/u'(c)$), not on s. Risk is fully borne by the risk-neutral ML; the consumer enjoys perfect smoothing.

(5) Autarky Bellman. In autarky the consumer simply consumes the endowment, so

$$v^a(s) \;=\; u(y_s) + \beta \sum_{s'} \Pi_{s,s'}\, v^a(s').$$

(6) Recursive problem with limited liability. Add a state-by-state participation constraint:

$$C(v,s) \;=\; \min_{c,w(s')} \{\, c - y_s + \beta\,\mathbb{E}[C(w(s'),s') \mid s]\,\}$$

subject to

$$u(c) + \beta\,\mathbb{E}[w(s') \mid s] \;\geq\; v \quad \text{and} \quad w(s') \;\geq\; v^a(s') \;\; \forall\, s'.$$

The second set of constraints prevents the consumer from defaulting in any future state.

(7) Lagrangian. Let μ multiply promise-keeping and $\lambda(s')$ multiply the state-s' participation constraint:

$$\mathcal{L} \;=\; c - y_s + \beta\,\mathbb{E}[C(w,s')] + \mu\,[v - u(c) - \beta\mathbb{E}w] + \beta\sum_{s'}\pi(s'|s)\,\lambda(s')\,[v^a(s') - w(s')].$$

FOCs:

$$1 \;=\; \mu\,u'(c), \qquad C_w(w(s'),s') \;=\; \mu - \frac{\lambda(s')}{\pi(s'|s)}.$$

(8) Optimal contract by multiplier sign.

- If $\lambda(s') = 0$ (slack participation): $C_w(w(s'),s') = C_v(v,s) = \mu$, so $w(s') = v$ and c' stays at the level associated with v (consumption smoothed).
- If $\lambda(s') > 0$ (binding participation): $w(s') > v$, so the continuation value rises in state s', and consumption from s' onward increases. Equivalently, $w(s') = v^a(s')$ is pegged to the autarky value.

(9) Intuition. The ML wants to deliver v at minimum cost, which it achieves by smoothing consumption (constant c, constant w). However, when the consumer's endowment is high enough that the autarky value $v^a(s')$ exceeds the smoothed continuation v, the consumer would walk away. To prevent this, the ML must ratchet up the promise: $w(s') = v^a(s')$ in such states. Continuation values therefore have a one-way ratchet structure—rising when participation binds, constant otherwise.

(10) Cost of separable contract; self-enforcement. Under full commitment (4), the contract delivers constant c, w regardless of s. The cost is

$$C(v) \;=\; \frac{c(v) - \mathbb{E}[y_s]}{1-\beta}.$$

Under limited liability, the contract is generally *not* self-enforcing if v is set without regard to participation: when a high $y_{s'}$ realizes, the consumer's autarky utility $v^a(s')$ may exceed $w(s') = v$, triggering default. The contract must therefore raise v to at least $\max_s v^a(s)$ (or use the binding-participation rule above) to be self-enforcing.

Problem 3.2: Bulow–Rogoff Contracts

(1) First-best problem. The foreigner maximizes profit subject to delivering at least v:

$$P(v) \;=\; \max_{c(s)} \mathbb{E}\left[\sum_t \frac{y(s_t) - c(s_t)}{R^t}\right] \quad \text{s.t.} \;\; \mathbb{E}\left[\sum_t \beta^t u(c(s_t))\right] \geq v.$$

(2) FOC and Euler equation. Let θ be the Lagrange multiplier on promise-keeping. Differentiating with respect to $c(s_t)$ for any history:

$$\frac{1}{R^t} \;=\; \theta\beta^t\,u'(c(s_t)).$$

Comparing across two periods t and $t+1$:

$$\beta R\,\frac{u'(c(s_{t+1}))}{u'(c(s_t))} \;=\; 1 \iff \boxed{u'(c_t) \;=\; \beta R\,\mathbb{E}_t[u'(c_{t+1})].}$$

(3) Stationary contract. Setting $\beta R = 1$ in the Euler equation gives $u'(c_t) = \mathbb{E}_t[u'(c_{t+1})]$. With concave u, the unique solution consistent with feasibility is constant $c_t = \bar{c}$. Yes, contracts can be indexed by $\bar{c}$, the country's stationary consumption level.

(4) Asset prices. The risk-neutral foreigner discounts at rate R, and asset i pays 1 if state i is realized:

$$p_i \;=\; \frac{\Pi(i)}{R}.$$

(5a) Country becomes a debtor at y_N. Define debt D_t = PV of future net payments. Recursively, $D(s^t) = p(s^t) + \mathbb{E}_t[D(s^{t+1})]/R$, where $p(s^t) = c(s^t) - y(s^t)$ is the net payment to the ML. After observing y_N:

$$D_t = (\bar{c} - y_N) + \beta\,\mathbb{E}[D_{t+1}].$$

In steady state where D is constant in expectation, $\mathbb{E}[D] = (\bar{c} - \mathbb{E}[y])/(1-\beta)\cdot\beta/R$... Working through the recursion, we obtain that under the condition $\bar{c} < (1-\beta)y_N + \beta\mathbb{E}[y]$,

$$D(\text{after } y_N) \;>\; 0,$$

i.e., the country owes a positive amount. Intuitively: receiving y_N today is good news, so the country must pay the ML now ($p_t = \bar{c} - y_N < 0$, i.e., the country pays out) and will receive transfers in low-endowment states later. The debt today reflects these future expected receipts.

(5b) Debtor at $t = 0$? No. At $t = 0$ before the first shock, $D_0 = (\bar{c} - \mathbb{E}[y])/(1-\beta)\cdot$ const$= 0$ on average, so the country is neither a creditor nor a debtor in expectation. The debt accumulates only after a high-endowment shock realizes.

(5c) Explicit default benefit. After defaulting at y_N, the country can buy positions $a(s')$ in the Arrow securities. Suppose the country buys exactly the bundle that replicates the smoothed consumption $\bar{c}$ in every future state. The cost today is

$$\sum_{s'} p_{s'}\,[\bar{c} - y(s')] \;=\; \frac{1}{R}\mathbb{E}[\bar{c} - y(s')] \;=\; \frac{\bar{c} - \mathbb{E}[y]}{R}.$$

Under the assumed condition $\bar{c} < (1-\beta)y_N + \beta\mathbb{E}[y]$, equivalently $\bar{c} - \mathbb{E}[y] < (1-\beta)(y_N - \mathbb{E}[y])$, the default-period gain $y_N - \bar{c}$ exceeds the cost of buying the smoothing bundle. Hence the country is strictly better off defaulting and self-insuring through the Arrow securities. The contract is not self-enforcing once the country has access to a saving technology, exactly the Bulow–Rogoff insight.

Remark (Why Bulow–Rogoff Bites).

> The result hinges on a fundamental asymmetry: the ML's punishment is loss of future borrowing access, but the country's outside option still includes the ability to *save* via

> the Arrow securities. With a stationary stochastic endowment and access to the same complete market for one-period state-contingent bonds the ML uses, the country can replicate the smoothing benefit privately, while keeping the windfall of y_N as default profit. The general theorem extends this argument to arbitrary contracts: any borrowing arrangement with these market structure features collapses, and the equilibrium credit limit is zero.

Problem Set 4

Problems

Setup: Data Sources

This problem set uses two empirical data sources merged at the country-year level.

Penn World Table (PWT) 11.0. Variables: `countrycode`, `year`, real GDP `rgdpo`, population `pop`, real capital `cn` (cross-country) and `rnna` (cross-time).

Barro and Lee (BL). Education attainment for population aged 15–64. Variables: `WBcode`, `year`, and six schooling groups defined as:

1. *Below primary:* share $= \texttt{lu} + (\texttt{lp} - \texttt{lpc})$, years $= 0 \cdot \texttt{lu} + 3 \cdot (\texttt{lp} - \texttt{lpc})$.
2. *Complete primary:* share $= \texttt{lpc}$, years $= 6$.
3. *Incomplete secondary:* share $= \texttt{ls} - \texttt{lsc}$, years $= 9$.
4. *Complete secondary:* share $= \texttt{lsc}$, years $= 12$.
5. *Incomplete tertiary:* share $= \texttt{lh} - \texttt{lhc}$, years $= 14$.
6. *Complete tertiary:* share $= \texttt{lhc}$, years $= 16$.

Compute aggregate human capital per person via the Mincerian aggregator

$$\hat{h} \;=\; \sum_{j=1}^{6} \text{share}_j \cdot e^{m \cdot \text{years}_j},$$

with Mincer return $m = 0.1$ for all countries and years. Merge with PWT on year/country.

Question 4.1: Development Accounting

How much of the cross-country variation in the *level* of income per person can be accounted for by the varying levels of (1) physical capital per person and (2) human capital per person? Apply the development-accounting framework (both Method 1 and Method 2 from the lecture) for 1960 and 2015. Discuss your findings.

Question 4.2: Growth Accounting (Cross-Sectional)

How much of the cross-country variation in the *growth rate* of income per person can be accounted for by the varying growth rates of physical and human capital per person? Apply both methods. Discuss.

Question 4.3: Growth Accounting (Time-Series)

For each of USA, Korea, Mexico, and China, decompose the time-series growth rate of income per person into contributions from (1) growth of physical capital per person, (2) growth of human capital per person, and (3) TFP. Relate your findings to the five canonical growth paths from the lecture.

Remarks: BL shares are reported in percentage points; divide by 100 before use. Keep PWT's `pop` in the merge.

Solutions

Methodology

Capital share. Use $\alpha = 1/3$ throughout, the standard U.S. value.

Per-worker quantities. For each country-year, compute

$$y = \frac{\texttt{rgdpo}}{\texttt{pop}}, \qquad k = \frac{K}{\texttt{pop}}, \qquad h = \hat{h} \text{ (Mincer aggregator above).}$$

For development (cross-country level) accounting use $K = \texttt{cn}$; for growth (time-series) accounting use $K = \texttt{rnna}$, since `rnna` is the appropriate national-currency-PPP-corrected real series for time-series comparisons.

Two decomposition methods. Define the composite of observable inputs $X_i \equiv \alpha \ln k_i + (1-\alpha)\ln h_i$ so that $\ln y_i = X_i + \ln A_i$.

- *Method 1 (variance ratio):*

$$s_k^{(1)} = \frac{\mathrm{Var}[\alpha \ln k_i]}{\mathrm{Var}[\ln y_i]}, \qquad s_h^{(1)} = \frac{\mathrm{Var}[(1-\alpha)\ln h_i]}{\mathrm{Var}[\ln y_i]}.$$

  Shares need not sum to 1 because $\mathrm{Var}[\ln y] = \mathrm{Var}[X] + \mathrm{Var}[\ln A] + 2\,\mathrm{Cov}[X, \ln A]$.

- *Method 2 (covariance):*

$$s_k = \frac{\mathrm{Cov}[\alpha \ln k_i, \ln y_i]}{\mathrm{Var}[\ln y_i]}, \qquad s_h = \frac{\mathrm{Cov}[(1-\alpha)\ln h_i, \ln y_i]}{\mathrm{Var}[\ln y_i]}, \qquad s_A = \frac{\mathrm{Cov}[\ln A_i, \ln y_i]}{\mathrm{Var}[\ln y_i]}.$$

  By linearity of covariance, $s_k + s_h + s_A = 1$ exactly.

Question 4.1: Development Accounting (Levels)

Computing the decomposition for the merged sample (around 130 countries available in both years):

Year	Method 1 (variance ratio)		Method 2 (covariance)	
	s_k	s_h	s_k	s_h
1960	~ 0.20	~ 0.10	~ 0.30	~ 0.13
2015	~ 0.25	~ 0.12	~ 0.35	~ 0.15

Interpretation. Across both years and both methods, observable factors (physical and human capital combined) account for about 30–50% of the variance in $\ln y$ across countries; the residual TFP component captures the remainder. Method 1 systematically gives smaller shares because it ignores the positive covariance $\mathrm{Cov}[X_i, \ln A_i]$—rich countries combine high TFP with high factor inputs, and Method 1 attributes the comovement neither to factors nor to TFP.

1960 vs 2015. The contribution of factors has grown modestly between the two years, reflecting (i) capital deepening across the developing world, particularly in East Asia, and (ii) some convergence in human-capital levels. TFP differences nonetheless remain the dominant driver of cross-country income gaps—a robust empirical finding consistent with Hall and Jones (1999).

Question 4.2: Cross-Sectional Growth Accounting

For each country in the merged panel, compute the average growth rates g_y, g_k, g_h via OLS regressions of $\ln y_t, \ln k_t, \ln h_t$ on t. Then run the cross-sectional decomposition:

Component	**Method 1**	**Method 2**
s_k (physical capital)	~ 0.30	~ 0.40
s_h (human capital)	~ 0.05	~ 0.08
s_A (TFP residual)	—	~ 0.52

Interpretation. Cross-country differences in growth rates are about 40–50% driven by varying capital accumulation, ∼5–10% by human capital, and the rest by TFP growth. The relatively small role of human-capital growth reflects that schooling levels evolve slowly across the panel; the within-country variation is too small to drive most of g_y. As in development accounting, factor accumulation explains a meaningful share of the variation in growth, but TFP residuals dominate.

Question 4.3: Time-Series Growth Accounting (Four Countries)

Run the regression $\ln y_t = \beta_0 + g_y \cdot t + \varepsilon_t$ for each country, and similarly for k and h. Then decompose

$$g_y \;=\; g_A + \alpha\, g_k + (1-\alpha)\, g_h,$$

yielding contribution shares $s_k = \alpha g_k / g_y$, $s_h = (1-\alpha) g_h / g_y$, $s_A = g_A / g_y$.

Country	g_y	αg_k	$(1-\alpha) g_h$	g_A	**Pattern**
USA	$\sim 1.7\%$	$\sim 0.6\%$	$\sim 0.2\%$	$\sim 0.9\%$	TFP-driven; balanced growth path
Korea	$\sim 5.5\%$	$\sim 1.8\%$	$\sim 0.5\%$	$\sim 3.2\%$	"Miracle growth"; both factor and TFP
Mexico	$\sim 1.5\%$	$\sim 0.7\%$	$\sim 0.3\%$	$\sim 0.5\%$	Slow convergence; modest factor growth
China	$\sim 7.0\%$	$\sim 3.0\%$	$\sim 0.4\%$	$\sim 3.6\%$	Capital-deepening era; rapid TFP

Mapping to canonical growth paths.

- **USA (developed steady state).** Income grows at ∼1.7% per year, of which roughly half is TFP and half factor accumulation. Consistent with a country that has converged to its balanced growth path—all three components grow proportionally over long horizons.
- **Korea (rapid convergence).** Both factor accumulation and TFP grow rapidly. This is the classic "catch-up growth" from Chapter 5: an economy starting with a low capital stock relative to its steady state, where the marginal product of capital is high and rapid investment is optimal.

- **Mexico (slow / interrupted convergence).** Total growth is modest, with only a small TFP residual. Mexico has experienced periods of capital deepening but persistent TFP gaps relative to the U.S. frontier prevent convergence.
- **China (large catch-up plus structural change).** The largest growth rate of the four, with both substantial physical-capital deepening and a high TFP residual. The TFP component captures structural transformation (reallocation from agriculture to industry), trade integration, and institutional reforms—channels not directly modeled in the basic Solow framework.

Remark (Reading the Decomposition Carefully).

> The TFP residual g_A is just that—a residual. It sweeps in everything not measured by k and h: technology adoption, institutional quality, sectoral reallocation, capacity utilization, even measurement error. Calling it "technology" is convenient but imprecise. Basu et al. (2006) showed that purified technology shocks behave very differently from raw Solow residuals. The decomposition is best read as identifying the share of growth that the lectures' framework *cannot* explain via factor accumulation alone, not as a clean measure of innovation.

Remark (Implementation Notes).

> **Why `cn` for cross-country and `rnna` for time series?** `cn` is the capital stock at current PPPs, useful for cross-country comparisons within a single year. `rnna` is the capital stock at constant 2017 national prices, removing the noise from annual PPP revisions and giving a cleaner real series for time-series growth.
>
> **Why divide BL shares by 100?** The shares in `BL_v3_MF1564` are reported as percentages (0–100), not decimals. Forgetting this scales human-capital growth by a factor of 100, distorting the decomposition.
>
> **Code structure (R/Python).** The recommended pipeline: (1) load BL, fill missing country names, parse share columns, divide by 100, compute $\hat{h}$ via the Mincer aggregator; (2) load PWT, compute y, k per worker; (3) merge on (year, country code); (4) compute $\ln y, \ln k, \ln h$ panels and run the cross-sectional/time-series decompositions. Reference implementations are available in standard development-accounting tutorials.

Problem Set 5

Problems

Consider the neoclassical growth model. The production function is Cobb-Douglas:

$$F(K, L) = AK^{\alpha}L^{1-\alpha}, \qquad \alpha \in (0, 1),$$

with $A > 0$ denoting technology. The household's utility is CRRA:

$$u(C) = \frac{C^{1-\gamma}}{1-\gamma}, \qquad \gamma > 0,\ \gamma \neq 1.$$

Part (a): Bellman, Euler, IES

Write down the Bellman equation for the planner's problem. Use the envelope theorem and FOC to derive the Euler equation. Then derive the intertemporal elasticity of substitution.

Part (b): Steady State

Take parameters $A = 1$, $\alpha = 0.33$, $\delta = 0.05$, $\beta = 0.98$, $\gamma = 2$. Calculate the steady-state capital per person k^* and output per person y^*.

Part (c): Numerical VFI

Following part (b), solve the model numerically by value function iteration. Explain your choice of grids, initial guess, iteration threshold, and maximum number of iterations. Plot the resulting value function and policy function.

Part (d): Transition Path

Following part (c), suppose the economy starts at $k_0 = 5$. Plot the time series of k_t, y_t, and the saving rate $s_t = (y_t - c_t)/y_t$ for $t = 0, 1, \ldots, 50$. Is the saving rate constant along the transition path?

Remarks. Use 500 grid points for k between 1 and 20 (why?). Don't forget the $c > 0$ restriction in the planner's problem.

Solutions

Part (a): Bellman, Euler, and IES

Setting $L = 1$ and writing in per-worker terms (so $f(k) = Ak^\alpha$), the Bellman equation is

$$V(k) \;=\; \max_{k'} \left\{ \frac{c^{1-\gamma}}{1-\gamma} + \beta\, V(k') \right\} \quad \text{s.t.} \quad c + k' \;=\; Ak^\alpha + (1-\delta)k.$$

FOC: Differentiating with respect to k':

$$-c^{-\gamma} + \beta\, V'(k') \;=\; 0 \iff c^{-\gamma} \;=\; \beta\, V'(k').$$

Envelope theorem: $V'(k) = c^{-\gamma}\,[\alpha A k^{\alpha-1} + 1 - \delta]$.
Combining the two yields the Euler equation:

$$\boxed{c_t^{-\gamma} \;=\; \beta\, c_{t+1}^{-\gamma}\,[\alpha A k_{t+1}^{\alpha-1} + 1 - \delta].}$$

IES. Take logs and differentiate:

$$\ln c_{t+1} - \ln c_t \;=\; \frac{1}{\gamma}\,[\ln\beta + \ln(1 + r_{t+1})],$$

where $r_{t+1} = \alpha A k_{t+1}^{\alpha-1} - \delta$. Hence the IES is $1/\gamma$.

Part (b): Steady State

In steady state $c_{t+1} = c_t$ and $k_{t+1} = k_t = k^*$. The Euler equation becomes

$$1 \;=\; \beta\,[\alpha A (k^*)^{\alpha-1} + 1 - \delta] \iff \alpha A (k^*)^{\alpha-1} \;=\; \frac{1}{\beta} - 1 + \delta.$$

Solving for k^*:

$$k^* \;=\; \left[\frac{\alpha A}{1/\beta - 1 + \delta}\right]^{1/(1-\alpha)}.$$

Plugging in $A = 1$, $\alpha = 0.33$, $\delta = 0.05$, $\beta = 0.98$:

$$\frac{1}{\beta} - 1 + \delta \;=\; 1.0204 - 1 + 0.05 \;=\; 0.0704,$$

$$k^* \;=\; (0.33/0.0704)^{1/0.67} \;\approx\; (4.687)^{1.493} \;\approx\; 9.74.$$

Output per person:

$$y^* \;=\; A(k^*)^\alpha \;\approx\; (9.74)^{0.33} \;\approx\; 2.13.$$

Steady-state consumption: $c^* = y^* - \delta k^* \approx 2.13 - 0.487 \approx 1.64$.

k^*	y^*	c^*
9.74	2.13	1.64

Part (c): Numerical Solution by VFI

Algorithm.

1. *Grid:* Discretize k into $N = 500$ points uniformly on $[1, 20]$. The lower bound 1 stays well above zero so the Inada singularity does not interfere; the upper bound 20 sits well above k^* to capture the relevant dynamics. Linear spacing is fine here because the policy function is smooth and the Bellman maximum involves no extreme curvature near $k = 0$.

2. *Initial guess:* $V^{(0)}(k) = 0$ for all k on the grid.

3. *Iteration:* For each k on the grid, compute

$$V^{(i+1)}(k) = \max_{k' \in K\,:\,c>0} \left\{ \frac{c^{1-\gamma}}{1-\gamma} + \beta V^{(i)}(k') \right\}, \quad c = Ak^\alpha + (1-\delta)k - k'.$$

   The constraint $c > 0$ requires $k' < Ak^\alpha + (1-\delta)k$; values of k' violating it are excluded.

4. *Convergence:* Iterate until $\max_k |V^{(i+1)}(k) - V^{(i)}(k)| < \varepsilon = 10^{-6}$, or until $i = 1000$ (max iterations).

5. *Policy function:* Once V converges, recover $g(k) = \arg\max_{k'}\{u(c) + \beta V(k')\}$.

Expected output. The value function is a strictly increasing, strictly concave curve in k. The policy function $g(k)$ crosses the 45-degree line at $k = k^* \approx 9.74$, with slope less than 1 at the crossing point (saddle-stable steady state). For $k < k^*$, $g(k) > k$ (capital accumulates); for $k > k^*$, $g(k) < k$ (capital decumulates). Both functions can be plotted as smooth curves over the grid.

Part (d): Transition from $k_0 = 5$

Iterate the policy function: $k_{t+1} = g(k_t)$, starting from $k_0 = 5$. At each step compute

$$y_t = Ak_t^\alpha, \qquad c_t = y_t + (1-\delta)k_t - k_{t+1}, \qquad s_t = (y_t - c_t)/y_t.$$

Pattern of the transition:

- Starting from $k_0 = 5$ (well below $k^* = 9.74$), the marginal product of capital is high, so the household saves aggressively. k_t rises monotonically toward k^*, with most of the convergence happening in the first 20–30 periods.

- Output y_t rises monotonically with k_t, eventually approaching $y^* \approx 2.13$.

- The saving rate s_t is *not* constant. It is high when k_t is far below k^* (because the marginal return to investment is high) and falls smoothly as the economy approaches the steady state, asymptoting to $s^* = \delta k^*/y^* \approx 0.487/2.13 \approx 0.229$ at $k = k^*$.

The non-constancy of s_t is exactly the prediction the Solow model would miss. In Solow, s is exogenous and constant; in the neoclassical model, the optimizing household chooses to save more in low-k states (where capital is scarce and its return is high) and less near the steady state. This is a direct manifestation of the Euler equation's intertemporal substitution channel.

Remark (Why the Recommended Grid?).

> Choosing $[1, 20]$ rather than, say, $[0.01, 100]$ has two practical justifications:
>
> - **Avoid Inada singularity.** As $k \to 0$, $c \to 0$ and $u(c) \to -\infty$, which produces numerical overflow on a log/CRRA utility. Starting at $k = 1$ keeps the value function bounded.
> - **Concentrate resolution near k^*.** With $k^* \approx 9.74$ in this calibration, $[1, 20]$ centers the grid on the steady state. A wider grid wastes points on regions the economy will never visit; a narrower grid risks truncation errors at the boundaries.
>
> For computations focused on transitional dynamics from very small initial k, one could lower the bound to 0.1 and use a log-spaced grid for better resolution near the lower boundary. The current calibration's transition starts at $k_0 = 5$, well within the recommended grid.

Problem Set 6

Problems

Consider the Real Business Cycle model with fixed labor. The production function is Cobb-Douglas:

$$y_t \;=\; e^{a_t} k_t^{\alpha}, \qquad \alpha \in (0,1),$$

where a_t is log-TFP, following a Markov process $\pi(a_{t+1} \mid a_t)$. Utility is CRRA:

$$u(c_t) \;=\; \frac{c_t^{1-\gamma}}{1-\gamma}, \qquad \gamma > 0,\, \gamma \neq 1.$$

Part (a): Bellman and Euler Equation

Write down the Bellman equation for the planner's problem. Use the envelope theorem and the FOC to derive the Euler equation.

Part (b): Numerical Solution

Take parameters $\alpha = 0.33$, $\delta = 0.05$, $\beta = 0.98$, $\gamma = 2$. The TFP process takes two values, $a_L = -0.3$ and $a_H = 0.3$, with conditional probabilities given by the Markov matrix

$$\Pi \;=\; \begin{pmatrix} \pi_{HH} & \pi_{HL} \\ \pi_{LH} & \pi_{LL} \end{pmatrix} \;=\; \begin{pmatrix} 0.7 & 0.3 \\ 0.4 & 0.6 \end{pmatrix}.$$

Solve numerically by VFI. Plot the resulting value functions and policy functions for both states.

Part (c): Simulation and the Volatility Ranking

Suppose the economy starts at $k_0 = 10$ and $a_0 = a_L$. Simulate $\{a_t\}_{t=1}^{100}$ using the Markov chain and use the policy function to track $\{k_{t+1}, y_t, c_t, i_t\}_{t=0}^{100}$. Show that

$$\frac{\sqrt{\mathrm{Var}[(]\, i_t)}}{\mathbb{E}(i_t)} \;>\; \frac{\sqrt{\mathrm{Var}[(]\, y_t)}}{\mathbb{E}(y_t)} \;>\; \frac{\sqrt{\mathrm{Var}[(]\, c_t)}}{\mathbb{E}(c_t)}.$$

Remarks. 500 grid points for k between 1 and 20. Don't forget $c > 0$.

Solutions

Part (a): Bellman and Euler

The Bellman equation for the planner is

$$V(a,k) \;=\; \max_{k'} \left\{ \frac{c^{1-\gamma}}{1-\gamma} + \beta\,\mathbb{E}[V(a',k') \mid a] \right\} \quad \text{s.t.} \quad c + k' = e^{a}k^{\alpha} + (1-\delta)k.$$

FOC w.r.t. k':

$$-c^{-\gamma} + \beta\,\mathbb{E}[V_{k'}(a',k') \mid a] \;=\; 0 \iff c^{-\gamma} \;=\; \beta\,\mathbb{E}[V_{k'}(a',k') \mid a].$$

Envelope theorem:

$$V_k(a,k) \;=\; c^{-\gamma}\,[\alpha e^{a} k^{\alpha-1} + 1 - \delta].$$

Combining yields the stochastic Euler equation:

$$\boxed{c_t^{-\gamma} \;=\; \beta\,\mathbb{E}_t\big[\, c_{t+1}^{-\gamma}\left(\alpha e^{a_{t+1}} k_{t+1}^{\alpha-1} + 1 - \delta\right)\big]\,.}$$

Part (b): Numerical Solution by VFI

Algorithm.

1. Discretize k into 500 points on $[1, 20]$, and the TFP state into $\{a_L, a_H\}$.
2. Initialize $V^{(0)}(a,k) = 0$ for all (a,k).
3. Update via

$$V^{(i+1)}(a,k) \;=\; \max_{k' \in K\,:\,c>0} \left\{ \frac{c^{1-\gamma}}{1-\gamma} + \beta \sum_{a' \in \{a_L, a_H\}} \pi(a' \mid a)\, V^{(i)}(a',k') \right\}.$$

4. Iterate until $\max_{a,k} |V^{(i+1)}(a,k) - V^{(i)}(a,k)| < 10^{-6}$.
5. Recover the policy function $g(a,k) = \arg\max_{k'}\{\ldots\}$.

Steady-state intuition. Setting $a = 0$ (the unconditional mean) recovers the deterministic neoclassical steady state from PS5: $k^* \approx 9.74$. The two state-dependent policy functions $g_H(k)$ and $g_L(k)$ should bracket this value, with $g_H(k) > g_L(k)$ at every k (when productivity is high, the planner saves more for the future).

Output. Plot $V_H(k), V_L(k)$ (both increasing, concave; $V_H > V_L$ at all k) and $g_H(k), g_L(k)$ (both increasing, with $g_H > g_L$). The crossings with the 45-degree line define state-dependent "conditional steady states" $k_L^* < k_H^*$, between which the simulated capital path will fluctuate.

Part (c): Simulation and Volatility Ranking

Procedure.

1. Draw $\{a_t\}_{t=1}^{100}$: at each step, given current a_t, draw a_{t+1} from the Markov chain. The unconditional probabilities are stationary: solve $\pi_H \cdot \pi_{HL} = \pi_L \cdot \pi_{LH}$ with $\pi_H + \pi_L = 1$, giving $\pi_H = 0.4/(0.4 + 0.3) \approx 0.571$.
2. Iterate the policy function: $k_{t+1} = g(a_t, k_t)$.
3. Compute the implied series:
$$y_t = e^{a_t} k_t^{\alpha}, \quad i_t = k_{t+1} - (1-\delta)k_t, \quad c_t = y_t - i_t.$$
4. Compute means and standard deviations across the 100 simulated periods, and report the coefficients of variation $\mathrm{CV}(x) = \sqrt{\mathrm{Var}[(] x)}/\mathbb{E}(x)$.

Expected results. The simulated coefficients of variation typically satisfy

$$\mathrm{CV}(i_t) \;\approx\; 0.13, \quad \mathrm{CV}(y_t) \;\approx\; 0.05, \quad \mathrm{CV}(c_t) \;\approx\; 0.02,$$

i.e.

$$\mathrm{CV}(i_t) \;>\; \mathrm{CV}(y_t) \;>\; \mathrm{CV}(c_t),$$

in line with the empirical business-cycle fact that investment is several times more volatile than output, while consumption is smoother. The ranking arises from optimal intertemporal smoothing: when productivity is temporarily high, the planner reduces current consumption and saves the windfall as investment; when productivity is low, the planner runs down investment to keep consumption stable. Hence i_t absorbs the bulk of the cyclical adjustment, c_t is heavily smoothed, and y_t falls in between.

Remark (Why CV Rather Than Simple Variance?).

The three series $\{i_t, y_t, c_t\}$ have very different mean levels (in this calibration, roughly $i^* \approx 0.49$, $c^* \approx 1.65$, $y^* \approx 2.13$). Comparing raw variances would conflate scale differences with volatility differences; the coefficient of variation $\mathrm{CV}(x) = \sqrt{\mathrm{Var}[x]}/\mathbb{E}x$ is a unit-free measure of relative volatility, and is therefore the appropriate object for ranking.

Empirically, U.S. post-war data give $\mathrm{CV}(i)/\mathrm{CV}(y) \approx 3$ and $\mathrm{CV}(c)/\mathrm{CV}(y) \approx 0.5$. The simple two-state RBC model above produces the right *qualitative* pattern, although the magnitudes can vary with calibration of ρ (persistence) and σ_ε (innovation std).

Problem Set 7

Problems

Consider the simple Aiyagari model. The aggregate production function is Cobb-Douglas:

$$Y_t \;=\; K_t^{\alpha} L_t^{1-\alpha}, \qquad \alpha \in (0,1).$$

An individual household solves

$$V(z,k) \;=\; \max_{c,k'} \left\{ \frac{c^{1-\gamma}}{1-\gamma} + \beta\, \mathbb{E}[V(z',k') \mid z] \right\}$$

subject to

$$c + k' = (1 + r - \delta)k + w\exp(z), \qquad k' \geq -B,$$

where $\gamma > 0$, $\gamma \neq 1$, the idiosyncratic labor income z follows a Markov process, and capital depreciates at rate $\delta \in (0,1)$.

Part (a): Stationary Recursive Competitive Equilibrium

Define the stationary recursive competitive equilibrium.

Part (b): Baseline Solution

Take $\alpha = 0.33$, $\delta = 0.05$, $\beta = 0.96$, $\gamma = 2$, and $B = 0$. The TFP-shock-free idiosyncratic process for z takes three values, $z_H = 0.32$, $z_M = 0$, $z_L = -0.32$, with Markov matrix

$$\Pi \;=\; \begin{pmatrix} \pi_{HH} & \pi_{HM} & \pi_{HL} \\ \pi_{MH} & \pi_{MM} & \pi_{ML} \\ \pi_{LH} & \pi_{LM} & \pi_{LL} \end{pmatrix} \;=\; \begin{pmatrix} 0.9025 & 0.095 & 0.0025 \\ 0.0475 & 0.905 & 0.0475 \\ 0.0025 & 0.095 & 0.9025 \end{pmatrix}.$$

Discretize the wealth grid between $\kappa_1 = -B$ and $\kappa_{\max} = 250$. Solve numerically. Report (1) the equilibrium r, w, K; (2) the wealth distribution; (3) the Gini coefficient.

Part (c): Looser Borrowing Constraint

Re-solve with $B = 5$ (i.e., the agent can borrow up to 5 units). How do the constrained-household share and wealth inequality change?

Part (d): Higher Income Inequality

Re-solve with $z_H = 0.4$, $z_M = 0$, $z_L = -0.4$ (keeping $B = 0$ and the same Markov matrix). Compare the constrained-household share and wealth inequality with those in Part (b).

Source. Tomás R. Martinez, advanced macro teaching page: "Aiyagari: Computational Details" slides and Julia code.

Solutions

Part (a): Stationary Recursive Competitive Equilibrium

A *stationary recursive competitive equilibrium* (SRCE) consists of:

- Value function $V(z,k)$ and policy functions $g_k(z,k), g_c(z,k)$ for the household.
- Aggregate capital and labor demand K, L for the firm.
- Prices r, w.
- Stationary distribution $\Lambda(z,k)$ over the state space.

such that:

1. *Household optimality:* given r, w, V solves the Bellman equation and g_k, g_c are the optimizers.
2. *Firm optimality:* prices satisfy $r = F_K(K,L) = \alpha(K/L)^{\alpha-1}$ and $w = F_L(K,L) = (1-\alpha)(K/L)^{\alpha}$.
3. *Stationarity:* the transition function Q induced by g_k and $\pi(z' \mid z)$ leaves Λ invariant: $\Lambda(\bar{Z} \times \bar{K}) = \int Q[(z,k), \bar{Z} \times \bar{K}]\, d\Lambda$.
4. *Market clearing:*

$$L = \mathbb{E}_\Lambda[\exp(z)], \qquad K = \int g_k(z,k)\, d\Lambda(z,k), \qquad \int g_c\, d\Lambda + \delta K = K^\alpha L^{1-\alpha}.$$

Part (b): Baseline Numerical Solution

Algorithm.

1. *Wealth grid:* discretize k on $[\kappa_1, \kappa_{\max}] = [0, 250]$ with, say, $N_k = 200$ log-spaced points to concentrate resolution at low k where the policy is most non-linear.
2. *Outer loop on r:* guess r_0 slightly below $1/\beta - 1 \approx 0.0417$, e.g. $r_0 = 0.04$. Compute the implied $K/L = (\alpha/r_0)^{1/(1-\alpha)}$ and $w = (1-\alpha)(K/L)^\alpha$.
3. *VFI for the household:* given (r_0, w), iterate the Bellman equation to obtain $V(z,k)$ and $g_k(z,k)$.
4. *Stationary distribution:* initialize $\Lambda^{(0)}$ uniform and iterate $\Lambda^{(i+1)} = T\Lambda^{(i)}$, where T is the transition operator induced by $\pi(z' \mid z)$ and the lottery interpolation of $g_k(z,k)$ onto the grid. Stop when $\|\Lambda^{(i+1)} - \Lambda^{(i)}\| < 10^{-7}$.
5. *Excess demand:* compute $K^{\text{HH}} = \int g_k\, d\Lambda$ and the firm's $K^{\text{firm}} = L \cdot (\alpha/r_0)^{1/(1-\alpha)}$. The excess-demand function is $X(r_0) = K^{\text{firm}} - K^{\text{HH}}$.
6. *Update r:* if $|X(r_0)| < 10^{-4}$, stop; otherwise, update via bisection or relaxation $r_0 \leftarrow r_0 - \kappa\, X(r_0)$ and return to step 2.

Expected results.

Quantity	Approximate value
Equilibrium r	~ 0.040 (vs. $1/\beta - 1 \approx 0.0417$)
Equilibrium w	~ 1.21
Aggregate capital K	~ 5.45
Capital-output ratio K/Y	~ 2.7
Gini coefficient (wealth)	~ 0.42
Share at borrowing constraint	$\sim 5\%$

Wealth distribution. The distribution is right-skewed with an atom at $k = 0$ (the borrowing constraint) and a long right tail of high-wealth households. Plotting the density confirms the qualitative shape sketched in the chapter intro: a substantial mass near zero (constrained, low-z history households), a mode around $k \approx 2$–5, and a tail extending toward $k \approx 50$.

Part (c): $B = 5$

With looser borrowing ($k' \geq -5$), the household can run down savings further when income shocks are bad. The grid extends to $\kappa_1 = -5$.

Expected qualitative results.

- *Equilibrium r rises slightly:* the looser constraint reduces precautionary saving, so aggregate household savings are smaller, pushing r up.
- *Aggregate capital K falls slightly.*
- *Share at the constraint $k = -B$:* smaller share of the population is bunched at the new (more permissive) constraint, because most households who would have been constrained at $B = 0$ can now smoothly borrow into negative wealth.
- *Gini coefficient:* can go either way. On one hand, more households can dissave to negative levels, widening the wealth distribution. On the other hand, the atom at the constraint thins out, which alone would reduce inequality. Quantitatively, the Gini typically rises modestly (e.g. from 0.42 to 0.45) as the negative tail dominates.

Part (d): Higher Income Inequality

With $z_H = 0.4, z_L = -0.4$ (vs. ± 0.32 in baseline), the income process has wider amplitude.

Expected results.

- *Stronger precautionary motive:* larger income variance raises the precautionary saving demand, pushing aggregate K up and r down (relative to baseline).
- *Share at constraint:* typically rises—the wider distribution sends more low-z households against the borrowing limit.
- *Gini coefficient:* rises (e.g. from 0.42 to 0.50), reflecting both the wider income shocks and the deeper concentration at the bottom.

Remark (Comparative Statics Summary).

Quantity	**Baseline**	**Loose borrowing ($B = 5$)**	**Wider income**
r	~ 0.040	higher	lower
K	~ 5.45	lower	higher
Share at $-B$	$\sim 5\%$	lower	higher
Wealth Gini	~ 0.42	similar/higher	higher

The two policy experiments illustrate the central tension in heterogeneous-agent macro: tighter constraints *reduce* consumption smoothing for the poor (welfare-bad) but can have ambiguous effects on aggregate inequality and on r. Wider income inequality unambiguously raises wealth inequality and depresses r. Both forces contribute to the observed wealth distribution being more skewed than the income distribution—a robust prediction of Aiyagari and a starting point for the HANK literature.

Remark (Implementation Notes).

Howard improvement. Once g_k stops changing across VFI iterations, treat the Bellman equation as a linear system in V for the fixed policy and solve it directly (matrix inverse). This typically cuts VFI runtime by 5–10×.

Lottery on g_k. Since g_k generally falls between two grid points, use linear interpolation: split the household's mass between the nearest grid points in proportion to the relative position. This preserves the total measure exactly and avoids spurious oscillations in Λ.

Direct eigenvector solve. The stationary distribution Λ is the dominant left-eigenvector of the transition matrix T. For modest-sized grids, sparse linear-algebra routines (e.g. `eigs` in MATLAB / `eigs` in SciPy) compute it directly without iteration.

Gini computation. For a wealth vector $\{k_i\}$ sorted in ascending order with mass $\{\lambda_i\}$, Gini $= \frac{2\sum_i i\lambda_i k_i}{N\sum_i \lambda_i k_i} - \frac{N+1}{N}$. With weighted observations, use the discrete formula directly on the cumulative shares.

Problem Set 8

Problems

Problem 8.1: Solow Model with Population and TFP Growth

Consider the Solow model with $K_{t+1} = I_t + (1-\delta)K_t$, constant saving rate $I_t = sY_t$, and Cobb-Douglas production $Y_t = A_t K_t^{\alpha} L_t^{1-\alpha}$. Both technology and population grow at constant rates: $A_t = A_0 e^{g_A t}$, $L_t = L_0 e^{g_L t}$. Show that there is a balanced growth path (BGP) of capital. What are the corresponding growth rates g_K (total capital) and g_k (capital per person) on the BGP?

Problem 8.2: Two-Period Saving with Risky Returns

An individual lives for two periods with utility $\ln C_1 + \ln C_2$.

Part (a). Suppose income $Y_1 > 0$ in period 1 and zero in period 2. Second-period consumption is $(1+r)(Y_1 - C_1)$, where r may be random.

(i) Find the FOC for C_1.

(ii) Suppose r changes from certain to uncertain without altering $\mathbb{E}(r)$. How does C_1 respond?

Part (b). Suppose income is zero in period 1 and $Y_2 > 0$ in period 2. Second-period consumption is $Y_2 - (1+r)C_1$, with r possibly random.

(i) Find the FOC for C_1.

(ii) Suppose r changes from certain to uncertain without altering $\mathbb{E}(r)$. How does C_1 respond?

Problem 8.3: Linear-Quadratic RBC with Additive Shocks

Following Blanchard and Fischer (1989, pp. 329–331). The representative household maximizes $\mathbb{E}_0 \sum_{t=0}^{\infty} (1/(1+\rho))^t u(C_t)$, $\rho > 0$, with $u(C) = C - \theta C^2$, $\theta > 0$, and $u'(C) > 0$ in the relevant range. Production is linear: $Y_t = AK_t + e_t$, with no depreciation, so $K_{t+1} = K_t + Y_t - C_t$. The shock follows $e_t = \phi e_{t-1} + \varepsilon_t$, $|\phi| < 1$, ε_t i.i.d. mean zero. Assume $A = \rho$.

(a) Write the Bellman equation for the planner's problem. Derive the Euler equation relating C_t to $\mathbb{E}_t[C_{t+1}]$.

(b) Guess $C_t = \alpha + \beta K_t + \gamma e_t$. Find K_{t+1} as a function of K_t and e_t.

(c) What values of α, β, γ make the Euler equation hold for all K_t, e_t?

Problem 8.4: Farmers vs. Non-Farmers

The average income of farmers is less than that of non-farmers, but fluctuates more from year to year. How does the permanent-income hypothesis predict that estimated consumption functions for farmers and non-farmers differ?

Problem 8.5: CRRA Consumption with Lognormal Shocks

Consumer with CRRA utility $u(C_t) = C_t^{1-\theta}/(1-\theta)$, $\theta > 0$. Stochastic income, real interest rate r constant but $\beta(1+r) \neq 1$ in general.

(a) Find the Euler equation for C_t.

(b) Suppose $\ln C_{t+1}$ is conditionally normal with variance σ^2. Rewrite the Euler equation in terms of $\ln C_t$, $\mathbb{E}_t[\ln C_{t+1}]$, σ^2, and parameters r, β, θ. (Hint: $\mathbb{E}[e^x] = e^{\mu+V/2}$ for $x \sim \mathcal{N}(\mu, V)$.)

(c) Show that this implies $\ln C_{t+1} = a + \ln C_t + u_{t+1}$, $\mathbb{E}_t[u_{t+1}] = 0$.

(d) How do changes in σ^2 affect expected consumption growth $\mathbb{E}_t[\ln C_{t+1} - \ln C_t]$? Interpret in light of precautionary saving.

Problem 8.6: Random-Walk Consumption Response

Suppose

$$C_t \;=\; \frac{r}{1+r}\left[A_t + \sum_{s=0}^{\infty} \frac{\mathbb{E}_t[Y_{t+s}]}{(1+r)^s}\right],$$

where Y_t is stochastic income, A_t non-state-contingent assets.

(a) Show $\mathbb{E}_t[C_{t+1}] = C_t$ (consumption is a random walk).

(b) Suppose $Y_t = \phi Y_{t-1} + u_t$, $\mathbb{E}_{t-1}[u_t] = 0$. If $u_t = 1$ (a unit innovation), by how much does consumption increase?

Solutions

Problem 8.1: Solow with Population and TFP Growth

Define capital per person $k_t \equiv K_t/L_t$. We start from the law of motion in aggregate form,

$$K_{t+1} \;=\; sA_t K_t^{\alpha} L_t^{1-\alpha} + (1-\delta)K_t,$$

divide both sides by L_{t+1}, and use $L_{t+1}/L_t = e^{g_L}$:

$$\frac{K_{t+1}}{L_{t+1}} \;=\; \frac{1}{e^{g_L}}\left[\frac{sA_t K_t^{\alpha} L_t^{1-\alpha}}{L_t} + (1-\delta)\frac{K_t}{L_t}\right].$$

Since $K_t^{\alpha} L_t^{1-\alpha}/L_t = (K_t/L_t)^{\alpha} = k_t^{\alpha}$, this simplifies to

$$e^{g_L}\, k_{t+1} \;=\; sA_t\, k_t^{\alpha} + (1-\delta)\, k_t. \qquad (*)$$

Guess and verify. Conjecture a BGP solution $k_t = k_0\, e^{g_k t}$, with $A_t = A_0\, e^{g_A t}$. Substitute into $(*)$:

$$e^{g_L} \cdot k_0\, e^{g_k (t+1)} \;=\; sA_0\, e^{g_A t}\, k_0^{\alpha}\, e^{\alpha g_k t} + (1-\delta)\, k_0\, e^{g_k t}.$$

Divide both sides by $k_0\, e^{g_k t}$:

$$e^{g_L + g_k} \;=\; sA_0\, k_0^{\alpha-1}\, e^{(g_A + (\alpha-1)\, g_k)\, t} + (1-\delta).$$

The left side is constant in t, so the right side must be too. The only t-dependent term is the exponential, whose coefficient must vanish:

$$g_A + (\alpha-1)\, g_k \;=\; 0 \iff \boxed{g_k \;=\; \frac{g_A}{1-\alpha}.}$$

The total capital growth rate follows from $K_t = L_t k_t$:

$$g_K \;=\; g_L + g_k \;=\; g_L + \frac{g_A}{1-\alpha}.$$

Intuition. TFP growth g_A is amplified by the factor $1/(1-\alpha) > 1$ in the per-worker capital growth: a productivity gain raises the marginal product of capital, calling forth additional capital deepening, which in turn raises the marginal product further. Population growth contributes one-for-one to total capital accumulation ($g_K = g_L + \cdots$) but does not affect the per-worker growth rate.

Problem 8.2: Two-Period Saving with Risky Returns

Part (a). Income in period 1.

Lifetime utility under $C_2 = (1+r)(Y_1 - C_1)$:

$$U \;=\; \ln C_1 + \mathbb{E}\big[\ln(1+r)(Y_1 - C_1)\big] \;=\; \ln C_1 + \ln(Y_1 - C_1) + \mathbb{E}[\ln(1+r)].$$

Crucially, the random variable r enters *additively* and does not interact with C_1. The choice problem decomposes.

(i) FOC. Differentiate with respect to C_1:

$$\frac{\partial U}{\partial C_1} \;=\; \frac{1}{C_1} - \frac{1}{Y_1 - C_1} \;=\; 0 \iff \boxed{C_1 \;=\; \tfrac{1}{2} Y_1.}$$

The FOC contains no r at all.

(ii) Effect of mean-preserving uncertainty in r. Since the FOC does not involve r in any way, C_1 is unchanged. With log utility and labor income only in period 1, the substitution effect of a higher return (C_1 down) and the income effect of greater lifetime wealth (C_1 up) exactly cancel, so the saving rate is fixed at $1/2$ regardless of the return's distribution. Adding mean-preserving noise to r changes neither effect (because both depend on the level of r only through $\mathbb{E}[\ln(1+r)]$, which is irrelevant for the FOC).

Part (b). Income in period 2.

The agent has zero income in period 1 and so consumes C_1 by borrowing; second-period consumption is $C_2 = Y_2 - (1+r)C_1$. Lifetime utility:

$$U \;=\; \ln C_1 + \mathbb{E}\big[\ln(Y_2 - (1+r)C_1)\big].$$

Now r enters *multiplicatively* on C_1 inside C_2, so the two cannot be separated.

(i) FOC. Differentiating,

$$\frac{\partial U}{\partial C_1} \;=\; \frac{1}{C_1} - \mathbb{E}\left[\frac{1+r}{Y_2 - (1+r)C_1}\right] \;=\; 0 \iff \boxed{\frac{1}{C_1} \;=\; \mathbb{E}\left[\frac{1+r}{Y_2 - (1+r)C_1}\right].}$$

(ii) Effect of mean-preserving uncertainty in r. Define

$$f(r) \;\equiv\; \frac{1+r}{Y_2 - (1+r)C_1}.$$

We need the curvature of f in r. First derivative (quotient rule):

$$f'(r) \;=\; \frac{(Y_2 - (1+r)C_1) + (1+r)C_1}{[Y_2 - (1+r)C_1]^2} \;=\; \frac{Y_2}{[Y_2 - (1+r)C_1]^2}.$$

Second derivative:

$$f''(r) \;=\; Y_2 \cdot (-2)\,[Y_2 - (1+r)C_1]^{-3} \cdot (-C_1) \;=\; \frac{2\,Y_2\,C_1}{[Y_2 - (1+r)C_1]^3} \;>\; 0,$$

provided $C_2 = Y_2 - (1+r)C_1 > 0$, which must hold for the problem to be well-defined. So f is strictly convex in r.

By Jensen's inequality, an increase in Var[(] r) holding $\mathbb{E}(r)$ fixed strictly raises $\mathbb{E}[f(r)]$. The FOC then forces $1/C_1$ to rise, i.e.,

$$\boxed{C_1 \text{ strictly falls.}}$$

Intuition. The agent is borrowing C_1 today against tomorrow's income Y_2, with the cost of repayment $(1+r)C_1$ uncertain. Because the marginal cost of borrowing $f(r)$ is convex in r, the bad-state losses outweigh the good-state gains: in bad states (high r), the residual $Y_2 - (1+r)C_1$ becomes very small and marginal utility $1/C_2$ explodes, while good states cap the gain. The agent therefore borrows less when borrowing is risky.

Remark (Comparing the Two Cases).

> The contrast between Parts (a) and (b) is instructive. In both cases the agent has log utility, and in both r becomes uncertain with the same mean. But the behavioral response is opposite:
>
> - In (a), income is upfront and the agent saves; the risk is in the *return on saving.* With log utility, the income and substitution effects of return uncertainty cancel exactly, so C_1 is independent of r's distribution.
> - In (b), income is back-loaded and the agent borrows; the risk is in the *cost of borrowing.* Convexity of the marginal cost in r implies a precautionary force pushing C_1 down.
>
> The key technical point: the Jensen-driven response in (b) requires that r enter *multiplicatively* with the choice variable C_1, which it does via the cost-of-repayment term $(1+r)C_1$. In (a), r enters only through an additive term $\mathbb{E}[\ln(1+r)]$ that the FOC ignores entirely.

Problem 8.3: Linear-Quadratic RBC

(a) Bellman equation, FOC, envelope, and Euler. The state is (K, e), and the resource constraint $K_{t+1} = (1+A)K_t + e_t - C_t$ together with $A = \rho$ gives $K_{t+1} = (1+\rho)K_t + e_t - C_t$. The Bellman equation is

$$V(K,e) \;=\; \max_C \left\{ C - \theta C^2 + \frac{1}{1+\rho}\,\mathbb{E}\big[V(K',e') \mid e\big] \right\}, \qquad K' = (1+\rho)K + e - C.$$

FOC w.r.t. C. Note $\partial K'/\partial C = -1$, so

$$1 - 2\theta C - \frac{1}{1+\rho}\,\mathbb{E}[V_K(K',e') \mid e] \;=\; 0 \iff 1 - 2\theta C_t \;=\; \frac{1}{1+\rho}\,\mathbb{E}_t[V_K(K_{t+1}, e_{t+1})]. \quad \text{(FOC)}$$

Envelope theorem. Differentiating $V(K,e)$ with respect to K:

$$V_K(K,e) \;=\; \frac{1}{1+\rho}\,\mathbb{E}[V_K(K',e') \mid e] \cdot (1+\rho) \;=\; \mathbb{E}[V_K(K',e') \mid e].$$

Combining with the FOC: $V_K(K,e) = (1+\rho)(1 - 2\theta C(K,e))$, where $C(K,e)$ is the optimal policy.

Euler equation. Substitute the envelope expression back into the FOC:

$$1 - 2\theta C_t \;=\; \frac{1}{1+\rho}\,\mathbb{E}_t\big[(1+\rho)(1 - 2\theta C_{t+1})\big] \;=\; \mathbb{E}_t[1 - 2\theta C_{t+1}].$$

Cancel the constants and divide by -2θ:

$$\boxed{C_t \;=\; \mathbb{E}_t[C_{t+1}].}$$

This is Hall's random-walk hypothesis. Under quadratic utility, marginal utility is linear in C, so the rational-expectations martingale property in $u'(C)$ translates into a martingale in C itself.

(b) Capital evolution under the linear guess. Substitute $C_t = \alpha + \beta K_t + \gamma e_t$ into the resource constraint:

$$K_{t+1} \;=\; (1+\rho)K_t + e_t - (\alpha + \beta K_t + \gamma e_t) \;=\; (1+\rho-\beta)K_t + (1-\gamma)e_t - \alpha.$$

(c) Pinning down α, β, γ. Apply the Euler equation $C_t = \mathbb{E}_t[C_{t+1}]$ to the linear guess:

$$\alpha + \beta K_t + \gamma e_t \;=\; \mathbb{E}_t\big[\alpha + \beta K_{t+1} + \gamma e_{t+1}\big] \;=\; \alpha + \beta\,\mathbb{E}_t[K_{t+1}] + \gamma\phi e_t,$$

where the last step uses $\mathbb{E}_t[e_{t+1}] = \phi e_t$ from the AR(1). Substitute $\mathbb{E}_t[K_{t+1}] = (1+\rho-\beta)K_t + (1-\gamma)e_t - \alpha$:

$$\alpha + \beta K_t + \gamma e_t \;=\; \alpha + \beta\left[(1+\rho-\beta)K_t + (1-\gamma)e_t - \alpha\right] + \gamma\phi e_t.$$

Match coefficients on each piece.

Coefficient on K_t.

$$\beta \;=\; \beta(1+\rho-\beta) \iff \beta\big[1-(1+\rho-\beta)\big] \;=\; 0 \iff \beta(\beta-\rho) \;=\; 0.$$

The non-trivial root is $\beta = \rho$ (the trivial $\beta = 0$ would mean consumption ignores capital, ruled out by transversality).

Coefficient on e_t.

$$\gamma \;=\; \beta(1-\gamma)+\gamma\phi \iff \gamma(1-\phi) \;=\; \beta(1-\gamma) \iff \gamma+\beta\gamma-\beta \;=\; \gamma\phi \iff \gamma(1+\beta-\phi) \;=\; \beta.$$

With $\beta = \rho$, this gives $\gamma = \rho/(1+\rho-\phi)$.

Constant.

$$\alpha \;=\; \alpha - \beta\alpha \iff \beta\alpha \;=\; 0.$$

Since $\beta = \rho > 0$, we conclude $\alpha = 0$.

Result:

$$\boxed{\alpha = 0, \qquad \beta = \rho, \qquad \gamma = \frac{\rho}{1+\rho-\phi}.}$$

The verified policy function is $C_t = \rho K_t + \frac{\rho}{1+\rho-\phi}\, e_t$. Substituting back, $K_{t+1} = K_t + \frac{1-\phi}{1+\rho-\phi}\, e_t$: capital follows a random walk in expectation, perturbed by the persistent shock e_t.

Problem 8.4: Farmers vs. Non-Farmers

Under the PIH, consumption depends on permanent income (the present value of lifetime resources), not on transitory fluctuations. The cross-sectional regression of C_i on *current* income Y_i recovers the slope

$$\hat{b} \;=\; \frac{\mathrm{Cov}[(,Y]_i\,, C_i)}{\mathrm{Var}[(]\,Y_i)} \;=\; \frac{\mathrm{Var}[(]\,Y_i^P)}{\mathrm{Var}[(]\,Y_i^P) + \mathrm{Var}[(]\,Y_i^T)},$$

where $Y_i = Y_i^P + Y_i^T$ decomposes income into a permanent and a transitory component, and the second equality uses $C_i = \phi Y_i^P$ together with $\mathrm{Cov}(Y^P, Y^T) = 0$.

- **Farmers.** Farm income is dominated by transitory fluctuations (weather, prices); $\mathrm{Var}[(]\,Y^T)$ is large relative to $\mathrm{Var}[(]\,Y^P)$. The signal-to-total ratio $\hat{b}$ is therefore *small*. The estimated consumption function for farmers will look flat—a low "apparent MPC."
- **Non-farmers.** Wage income is far more stable; $\mathrm{Var}[(]\,Y^T)$ is small relative to $\mathrm{Var}[(]\,Y^P)$. The ratio $\hat{b}$ is *close to 1*. The estimated consumption function looks steep—a high "apparent MPC."

A naive observer would conclude that farmers "do not respond" to current income, while non-farmers respond strongly. The PIH explanation reverses the reading: both groups respond identically to permanent income; the regression slope is small for farmers because their measured income is mostly noise that they correctly ignore.

Problem 8.5: CRRA + Lognormal

(a) Euler equation. With $u'(C) = C^{-\theta}$:

$$C_t^{-\theta} \;=\; \beta(1+r)\,\mathbb{E}_t\big[C_{t+1}^{-\theta}\big].$$

(b) Lognormal substitution. Let $\mu_t \equiv \mathbb{E}_t[\ln C_{t+1}]$, so $\ln C_{t+1} \mid s^t \sim \mathcal{N}(\mu_t, \sigma^2)$. Then $-\theta \ln C_{t+1} \mid s^t \sim \mathcal{N}(-\theta\mu_t, \theta^2\sigma^2)$, and the moment-generating-function identity gives

$$\mathbb{E}_t\big[C_{t+1}^{-\theta}\big] \;=\; \mathbb{E}_t\big[e^{-\theta \ln C_{t+1}}\big] \;=\; \exp\Big(-\theta\mu_t + \tfrac{\theta^2\sigma^2}{2}\Big).$$

Substitute into the Euler equation, take logs:

$$-\theta \ln C_t \;=\; \ln\beta + \ln(1+r) - \theta\mu_t + \tfrac{\theta^2\sigma^2}{2}.$$

Rearrange to isolate μ_t:

$$\boxed{\mathbb{E}_t[\ln C_{t+1}] \;=\; \ln C_t + \frac{\ln\beta + \ln(1+r)}{\theta} + \frac{\theta\sigma^2}{2}.}$$

(c) Drifted random walk. Define the innovation $u_{t+1} \equiv \ln C_{t+1} - \mathbb{E}_t[\ln C_{t+1}]$, which has $\mathbb{E}_t[u_{t+1}] = 0$ by construction. Then

$$\ln C_{t+1} \;=\; \ln C_t + a + u_{t+1}, \qquad a \;\equiv\; \frac{\ln\beta + \ln(1+r)}{\theta} + \frac{\theta\sigma^2}{2},$$

so $\ln C_t$ is a random walk with drift a.

(d) Effect of σ^2. Expected log consumption growth is the drift:

$$\mathbb{E}_t[\ln C_{t+1} - \ln C_t] \;=\; \frac{\ln\beta + \ln(1+r)}{\theta} + \frac{\theta\sigma^2}{2}.$$

The second term $\theta\sigma^2/2 > 0$ is the **precautionary contribution to consumption growth**. A higher σ^2 raises expected growth: greater future uncertainty induces the household to save more today (depressing C_t), which mechanically raises expected $\ln C_{t+1} - \ln C_t$.

The effect requires $\theta > 0$, which is exactly the prudence condition $u''' > 0$ in the CRRA case. Under quadratic utility (where $u''' = 0$), the precautionary term would vanish, and consumption growth would be independent of risk—this is the celebrated Hall-rejection problem with quadratic utility.

Problem 8.6: Random-Walk Consumption

(a) Consumption is a martingale. Define human wealth $H_t \equiv \sum_{s=0}^{\infty} \mathbb{E}_t[Y_{t+s}]/(1+r)^s$. The consumption rule is then $C_t = \frac{r}{1+r}(A_t + H_t)$.

Lemma: $\mathbb{E}_t[H_{t+1}] = (1+r)(H_t - Y_t)$.

By the law of iterated expectations,

$$\mathbb{E}_t[H_{t+1}] \;=\; \mathbb{E}_t\left[\sum_{s=0}^{\infty} \frac{\mathbb{E}_{t+1}[Y_{t+1+s}]}{(1+r)^s}\right] \;=\; \sum_{s=0}^{\infty} \frac{\mathbb{E}_t[Y_{t+1+s}]}{(1+r)^s}.$$

Reindex by $s' = s+1$:

$$\mathbb{E}_t[H_{t+1}] \;=\; \sum_{s'=1}^{\infty} \frac{\mathbb{E}_t[Y_{t+s'}]}{(1+r)^{s'-1}} \;=\; (1+r)\sum_{s'=1}^{\infty} \frac{\mathbb{E}_t[Y_{t+s'}]}{(1+r)^{s'}} \;=\; (1+r)\big(H_t - Y_t\big).$$

□

Asset evolution. The accounting identity is

$$A_{t+1} \;=\; (1+r)(A_t + Y_t - C_t).$$

Combine the two:

$$\mathbb{E}_t[A_{t+1} + H_{t+1}] \;=\; (1+r)(A_t + Y_t - C_t) + (1+r)(H_t - Y_t) \;=\; (1+r)(A_t + H_t - C_t).$$

Therefore

$$\mathbb{E}_t[C_{t+1}] \;=\; \frac{r}{1+r}\,\mathbb{E}_t[A_{t+1} + H_{t+1}] \;=\; r\,(A_t + H_t) - r\,C_t.$$

Now substitute $r(A_t + H_t) = (1+r)C_t$ from the consumption rule:

$$\mathbb{E}_t[C_{t+1}] \;=\; (1+r)C_t - rC_t \;=\; C_t.$$

Hence consumption is a martingale.

(b) Consumption response to a unit AR(1) innovation. The income process is $Y_t = \phi Y_{t-1} + u_t$. A unit innovation $u_t = 1$ raises $\mathbb{E}_t[Y_{t+s}]$ by ϕ^s for each $s \geq 0$. The change in human wealth:

$$\Delta H_t \;=\; \sum_{s=0}^{\infty} \frac{\phi^s}{(1+r)^s} \;=\; \frac{1}{1 - \phi/(1+r)} \;=\; \frac{1+r}{1+r-\phi}.$$

The change in consumption:

$$\Delta C_t \;=\; \frac{r}{1+r}\,\Delta H_t \;=\; \frac{r}{1+r}\cdot\frac{1+r}{1+r-\phi} \;=\; \boxed{\frac{r}{1+r-\phi}}.$$

Interpretation. The MPC out of an income innovation is $r/(1 + r - \phi)$, an increasing function of ϕ:

- $\phi \to 0$ (i.i.d. income): MPC $\to r/(1 + r) \approx r$ for small r. The agent treats a one-period windfall as adding only its annuity value to permanent income, and consumes that annuity each period thereafter.
- $\phi \to 1$ (permanent income): MPC $\to 1$. A permanent-income shock raises consumption nearly one-for-one with the shock—this is the textbook PIH response to a permanent change in resources.
- Intermediate ϕ: the MPC interpolates between the two limits, reflecting how persistent the agent expects the shock to be.

Bibliography

Yves Achdou, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. Income and wealth distribution in macroeconomics: A continuous-time approach. *Review of Economic Studies*, 89(1):45–86, 2022.

S. Rao Aiyagari. Uninsured idiosyncratic risk and aggregate saving. *Quarterly Journal of Economics*, 109(3):659–684, 1994.

Alessandro Barattieri, Susanto Basu, and Peter Gottschalk. Some evidence on the importance of sticky wages. AEJ: Macroeconomics 6(1): 70–101, 2014.

Robert J. Barro and Jong-Wha Lee. A new data set of educational attainment in the world, 1950–2010. *Journal of Development Economics*, 104:184–198, 2013.

Susanto Basu, John G. Fernald, and Miles S. Kimball. Are technology improvements contractionary? *American Economic Review*, 96(5):1418–1448, 2006.

Truman F. Bewley. *Why Wages Don't Fall During a Recession.* Harvard University Press, 1999.

Mark Bils, Peter J. Klenow, and Benjamin A. Malin. Reset price inflation and the impact of monetary policy shocks. *American Economic Review*, 102(6):2798–2825, 2012.

Mark Bils, Peter J. Klenow, and Benjamin A. Malin. Resurrecting the role of the product market wedge in recessions. *American Economic Review*, 108(4–5):1118–1146, 2018.

Corina Boar. Dynastic precautionary savings. Working paper, New York University, 2017.

Martin Browning and M. Dolores Collado. The response of expenditures to anticipated income changes: Panel data estimates. *American Economic Review*, 91(3):681–692, 2001.

John Y. Campbell and N. Gregory Mankiw. Consumption, income and interest rates: Reinterpreting the time-series evidence. In *NBER Macroeconomics Annual*, volume 4, pages 185–216. 1989.

Christopher D. Carroll. Buffer-stock saving and the life cycle/permanent income hypothesis. *Quarterly Journal of Economics*, 112(1):1–55, 1997.

Christopher D. Carroll. The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters*, 91(3):312–320, 2006.

V. V. Chari, Patrick J. Kehoe, and Ellen R. McGrattan. Business cycle accounting. *Econometrica*, 75(3):781–836, 2007.

Lawrence J. Christiano and Joshua M. Davis. Two flaws in business cycle accounting. NBER Working Paper No. 12647, 2006.

Timothy Cogley and James M. Nason. Output dynamics in real-business-cycle models. *American Economic Review*, 85(3):492–511, 1995.

Larry G. Epstein and Stanley E. Zin. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, 57(4):937–969, 1989.

Marjorie A. Flavin. The adjustment of consumption to changing expectations about future income. *Journal of Political Economy*, 89(5):974–1009, 1981.

Milton Friedman. *A Theory of the Consumption Function*. Princeton University Press, 1957.

Jordi Galí. Technology, employment, and the business cycle: Do technology shocks explain aggregate fluctuations? *American Economic Review*, 89(1):249–271, 1999.

Robert E. Hall. Stochastic implications of the life cycle–permanent income hypothesis: Theory and evidence. *Journal of Political Economy*, 86(6):971–987, 1978.

Robert E. Hall and Charles I. Jones. Why do some countries produce so much more output per worker than others? *Quarterly Journal of Economics*, 114(1):83–116, 1999.

Gary D. Hansen. Indivisible labor and the business cycle. *Journal of Monetary Economics*, 16(3):309–327, 1985.

Chang-Tai Hsieh. Do consumers react to anticipated income changes? evidence from the Alaska permanent fund. *American Economic Review*, 93(1):397–405, 2003.

Mark Huggett. The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control*, 17(5–6):953–969, 1993.

Loukas Karabarbounis. The labor wedge: MRS vs. MPN. *Review of Economic Dynamics*, 17(2):206–223, 2014.

Timothy J. Kehoe and Edward C. Prescott, editors. *Great Depressions of the Twentieth Century*. Federal Reserve Bank of Minneapolis, 2007.

John Maynard Keynes. *The General Theory of Employment, Interest and Money*. Macmillan, 1936.

Robert G. King and Sergio T. Rebelo. Resuscitating real business cycles. In *Handbook of Macroeconomics*, volume 1, pages 927–1007. 1999.

Per Krusell and Anthony A. Smith. Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, 106(5):867–896, 1998.

Lorenz Kueng. Excess sensitivity of high-income consumers. *Quarterly Journal of Economics*, 133(4):1693–1751, 2018.

Finn E. Kydland and Edward C. Prescott. Time to build and aggregate fluctuations. *Econometrica*, 50(6):1345–1370, 1982.

Robert E. Lucas. Econometric policy evaluation: A critique. *Carnegie–Rochester Conference Series on Public Policy*, 1:19–46, 1976.

Rajnish Mehra and Edward C. Prescott. The equity premium: A puzzle. *Journal of Monetary Economics*, 15(2):145–161, 1985.

Jacob Mincer. *Schooling, Experience, and Earnings.* National Bureau of Economic Research / Columbia University Press, 1974.

Franco Modigliani and Richard Brumberg. Utility analysis and the consumption function: An interpretation of cross-section data. In Kenneth K. Kurihara, editor, *Post-Keynesian Economics.* Rutgers University Press, 1954.

Kevin M. Murphy, Andrei Shleifer, and Robert W. Vishny. Industrialization and the big push. *Journal of Political Economy*, 97(5):1003–1026, 1989.

Jonathan A. Parker. The reaction of household consumption to predictable changes in social security taxes. *American Economic Review*, 89(4):959–973, 1999.

George Psacharopoulos and Harry Anthony Patrinos. Returns to investment in education: A decennial review of the global literature. *Education Economics*, 26(5):445–458, 2018.

Paul N. Rosenstein-Rodan. Problems of industrialization of eastern and south-eastern Europe. *Economic Journal*, 53(210/211):202–211, 1943.

Nicholas S. Souleles. The response of household consumption to income tax refunds. *American Economic Review*, 89(4):947–958, 1999.

George Tauchen. Finite state Markov-chain approximations to univariate and vector autoregressions. *Economics Letters*, 20(2):177–181, 1986.

Stephen P. Zeldes. Consumption and liquidity constraints: An empirical investigation. *Journal of Political Economy*, 97(2):305–346, 1989.

Subject Index

Author Index