Preface

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

Audience and Prerequisites

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

Structure of the Book

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

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

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

Pedagogical Conventions

Several typographic conventions recur throughout the text.

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

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

Reading Paths

Readers do not have to proceed linearly.

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

Acknowledgments

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

Rui Zhou, Spring 2026

Notation

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

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

Symbols used throughout the book.

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

(continued on next page)

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

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

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

Part I

Heterogeneous Agents in Complete and Incomplete Markets

Lectures by Maria-Jose Carreras-Valle

Part II

Growth, Business Cycles, and Quantitative Macroeconomics

Lectures by Kai-Jie Wu

Chapter 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 alloca-

> tions (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.

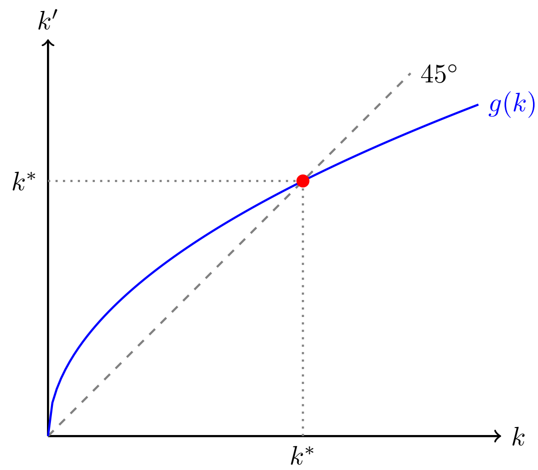

This phase diagram illustrates the dynamic evolution of capital via the policy function $k' = g(k)$. The horizontal axis represents the current capital stock k, and the vertical axis represents tomorrow's capital stock k'. The dashed 45° line represents the locus of points where $k' = k$ (i.e., capital 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).

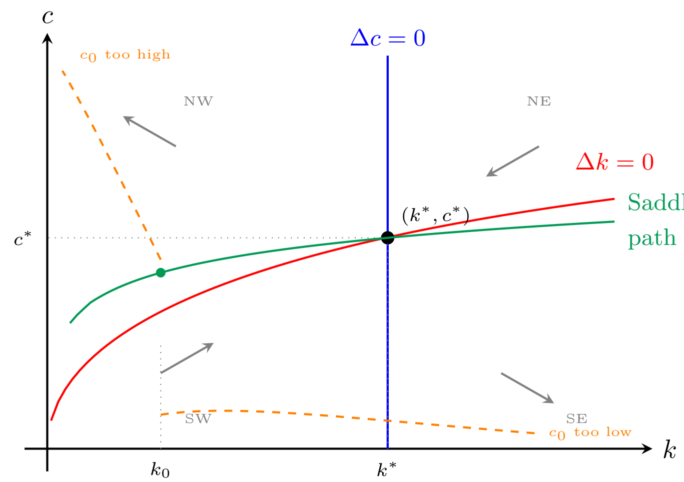

Three features of the picture are worth flagging.

- **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.

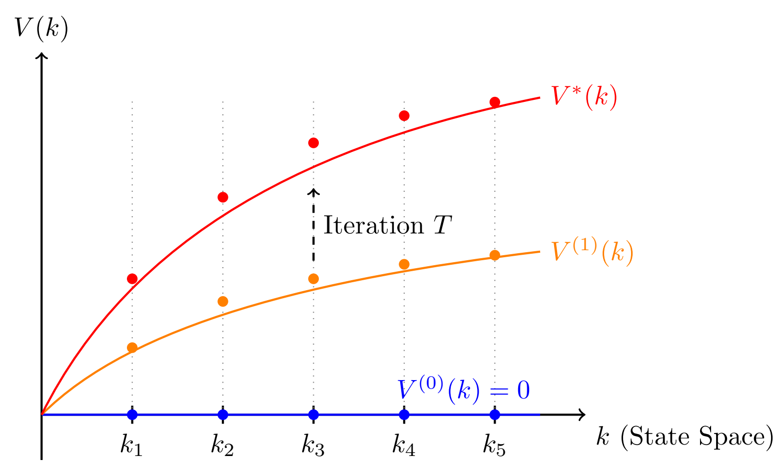

Remark (Continuous Choice, Interpolation, and the Curse of Dimensionality).

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.

Part III

Problem Sets and Solutions

Subject Index

Author Index