Preface

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

Audience and Prerequisites

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

Structure of the Book

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

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

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

Pedagogical Conventions

Several typographic conventions recur throughout the text.

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

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

Reading Paths

Readers do not have to proceed linearly.

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

Acknowledgments

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

Rui Zhou, Spring 2026

Notation

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

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

Symbols used throughout the book.

Symbol	Meaning
Preferences and discounting	
$u(\cdot)$	Period utility function; $u' > 0$, $u'' < 0$, satisfying Inada conditions where needed.
β	Time discount factor; $\beta \in (0,1)$.
σ	Coefficient of relative risk aversion under CRRA utility; the inverse $1/\sigma$ is the intertemporal elasticity of substitution.
γ	Coefficient of *absolute* risk aversion under CARA utility (Ch. 2 only).
$\mathbb{E}_t[\cdot]$	Expectation conditional on history s^t.
Stochastic environment	
s_t, s^t	Date-t state; history through t.
$\pi(s^t)$	Unconditional probability of history s^t; $\pi(s^\tau \mid s^t)$ 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

Part II

Growth, Business Cycles, and Quantitative Macroeconomics

Lectures by Kai-Jie Wu

Chapter 7

Neoclassical Growth vs. Data

Remark (Notation in This Chapter).

Symbol	Meaning
$g_y^{\text{1960-1990}}$	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.

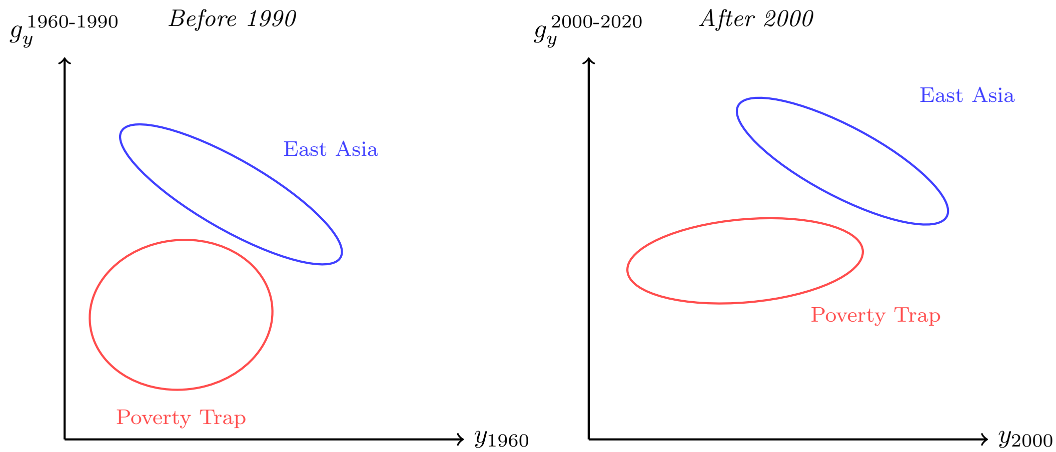

Two features of the data stand out. First, prior to 1990, the cross-country distribution of (initial income, growth rate) pairs is strikingly *bimodal*. A cluster of East Asian economies—South Korea, Taiwan, Singapore, Hong Kong—achieved rapid catch-up growth of precisely the kind the Solow model predicts, 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:

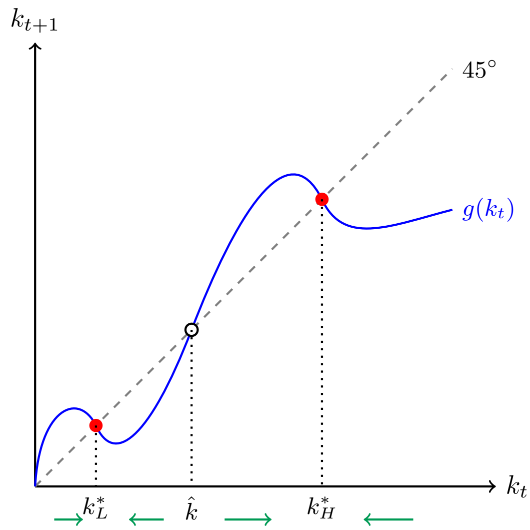

The three intersections yield two stable and one unstable steady state:

- 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 $k_H^* > k_L^*$.

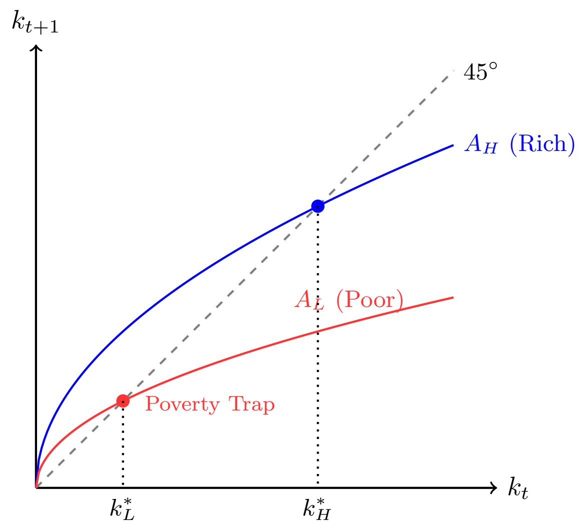

Under this mechanism, the poverty trap is not a trap in the traditional sense of a basin of attraction with a tipping point. The poor economy is simply converging to a *different*, structurally lower steady state. No matter how long we wait, it will not catch up to the rich economy—unless its structural 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 **?** in a model of industrialization with increasing returns to scale, building on **?**'s (**?**) 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.

Part III

Problem Sets and Solutions

Subject Index

Author Index