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	
Var, 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 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).

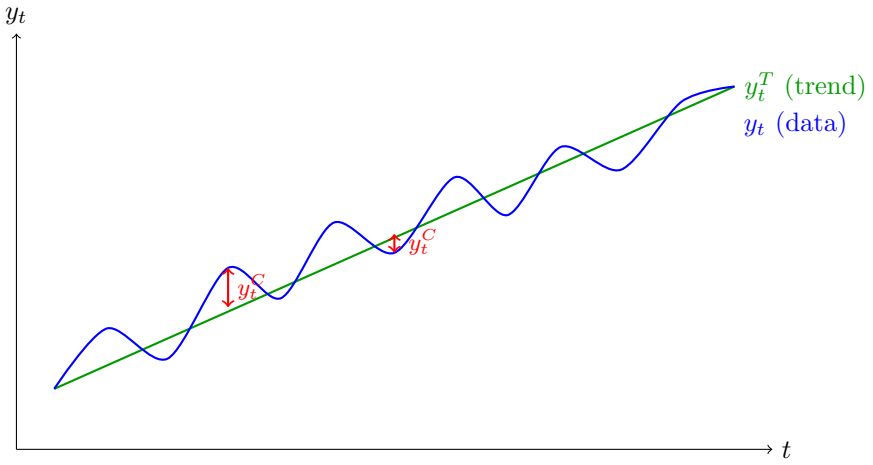

Remark (The Trend-Cycle Decomposition is Not Unique).

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.

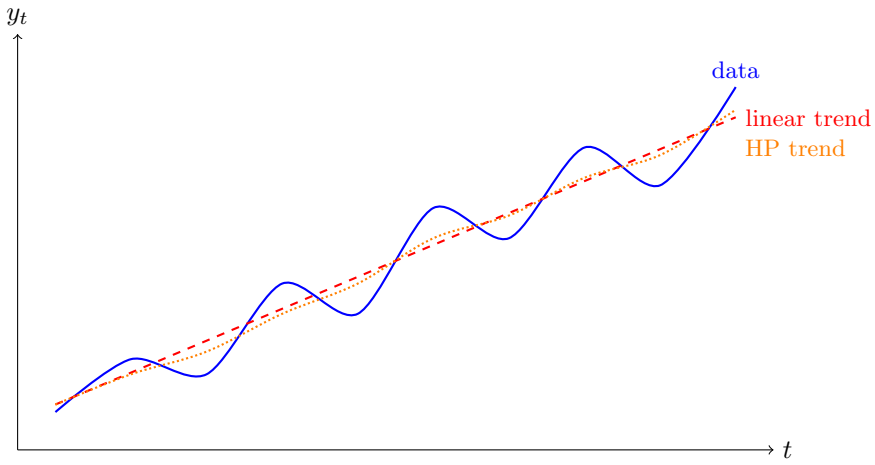

A **recession** is informally defined as two or more consecutive quarters in which *real GDP falls relative to the previous quarter* (i.e., negative quarter-over-quarter growth), or equivalently, a sustained rise in the unemployment rate. This is a statement about the *level* of GDP, not about its deviation 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 **?**. 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 (**?**, **?**).

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.

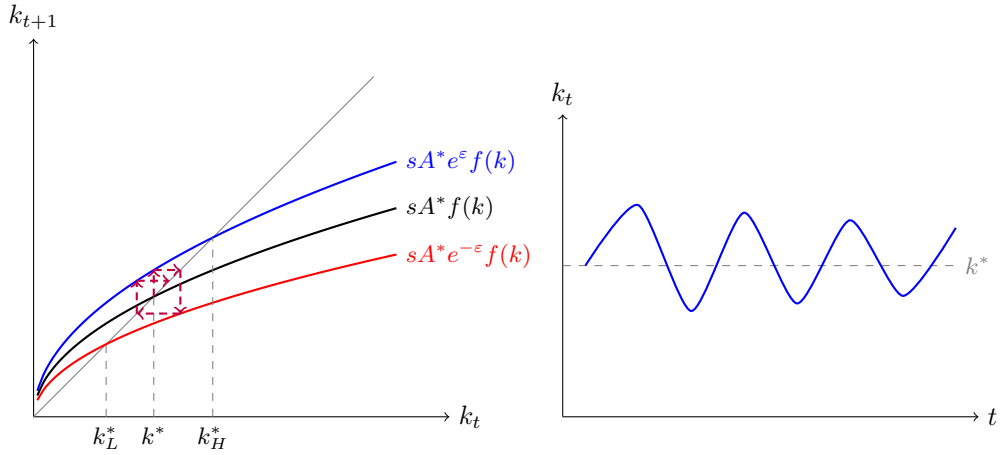

However, the Solow model imposes a **constant saving rate**: $I_t = sY_t$. This immediately implies

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

Taking variances of both sides:

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

so $\mathrm{Var}[\ln I_t] = \mathrm{Var}[\ln Y_t]$. 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 **?**. 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 $k_{t+1} = g(a_t, k_t)$.

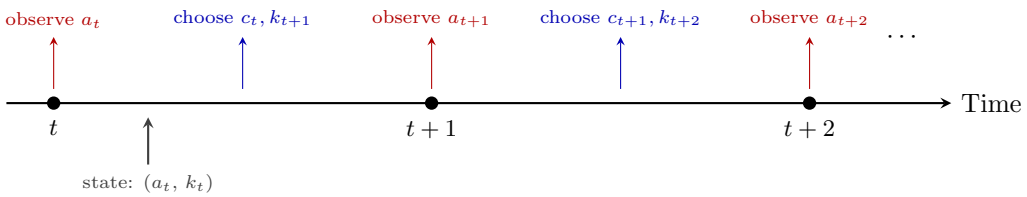

8.4.2 The Planner's Problem

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:

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

> 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 $g_L^n \equiv g(a_L, k^n)$.

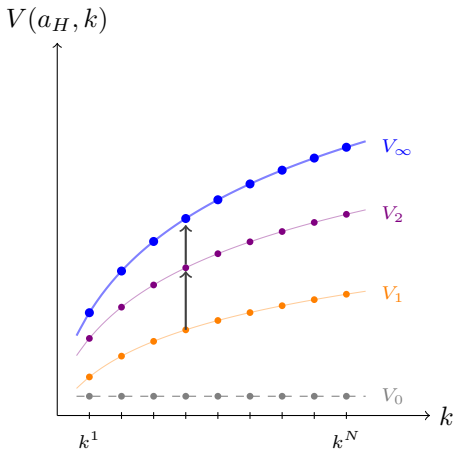

Value Function Iteration

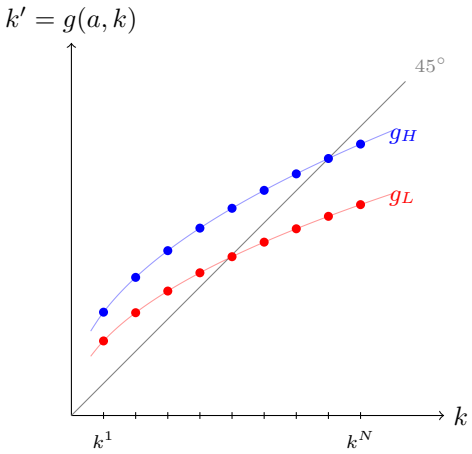

Policy Functions

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:

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

and analogously for the low state:

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

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:

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

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**.

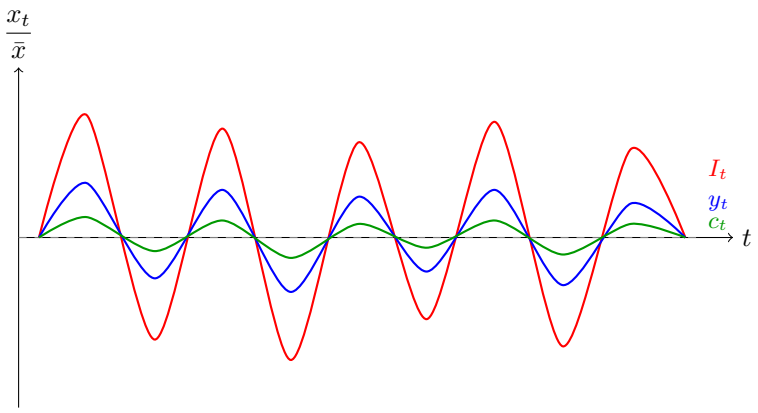

Impulse Response: Anatomy of a Single Shock

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)**.

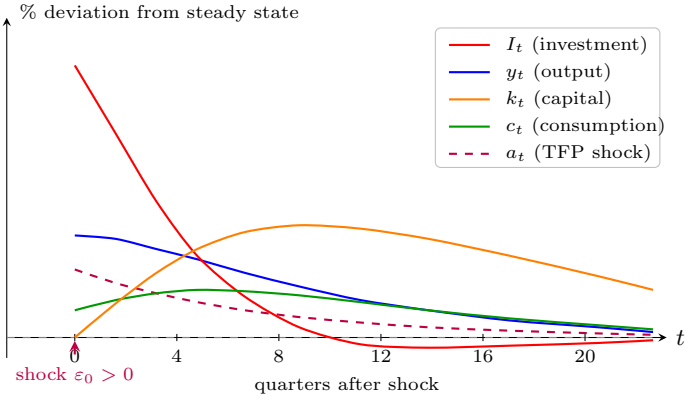

Three features of this picture are worth flagging:

- **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 (**?**, **?**) 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 dispro-

portionately 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:

$$\vec{a} = [a^1, a^2, \ldots, a^M],$$

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 **?**. The idea is to choose the grid points $\{a^1, \ldots, a^M\}$ 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:

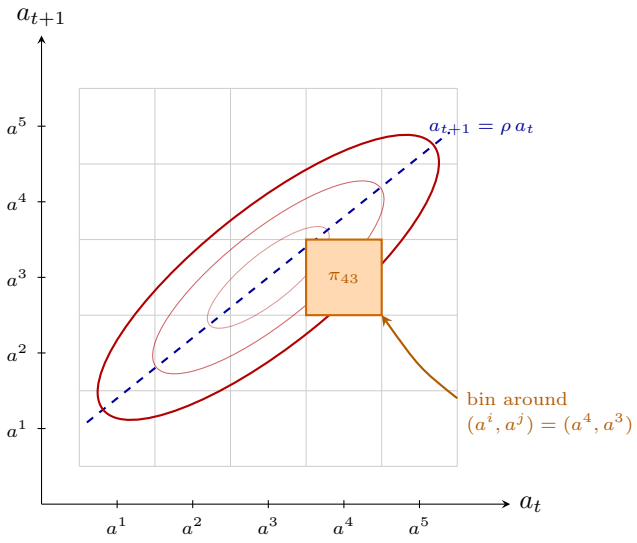

The red ellipses are level sets of the joint density $f(a_t, a_{t+1})$; the dashed line is the conditional mean $\mathbb{E}(a_{t+1} \mid a_t) = \rho a_t$. The Tauchen probability π_{ij} is the conditional mass of the orange bin centered at (a^i, a^j) given $a_t = a^i$—i.e., the integral of the conditional 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:

$$V_{m,\mathtt{nI}+1}^n = \max_{n' \in \{1,\ldots,N\}} \left\{ u\left[e^{a^m} f(k^n) + (1-\delta)k^n - k^{n'}\right] + \beta \sum_{j=1}^{M} \pi_{mj}\, V_{j,\mathtt{nI}}^{n'} \right\}.$$

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 $a^m < a^{m'}$.

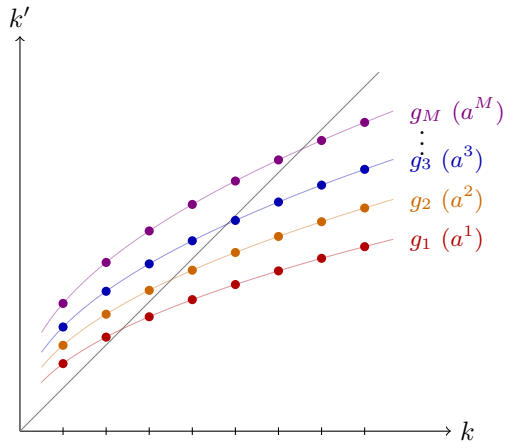

Step 5: Simulate

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.

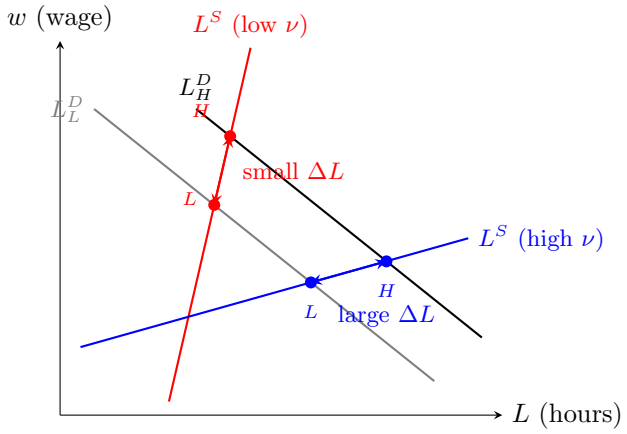

As the diagram shows, a steep labor supply curve (low ν, red) implies that the same demand shift produces only a small change in equilibrium hours—most of the adjustment

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:

$$\hat{\nu}_{\text{micro}} \ll 1, \quad \text{typically 0.1–0.2.}$$

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 (? and Indivisible Labor).

A classic theoretical device for closing the gap is **?**'s (**?**) **indivisible labor** model. Suppose individual workers face a fixed cost of going to work (commuting, childcare, 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.

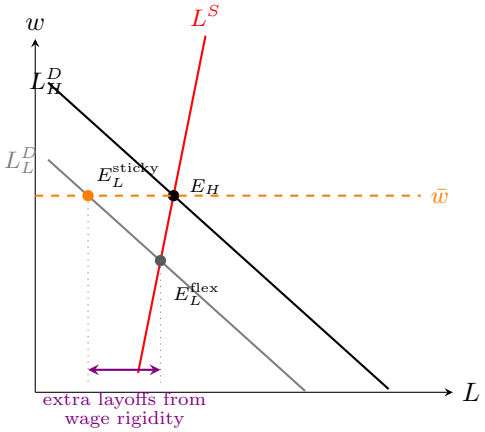

Under flexible wages, a negative shock moves the equilibrium from E_H to E_L^{flex}: both wages and hours fall modestly. Under sticky wages, the wage stays at $\bar{w}$, so the firm slides along its demand curve to E_L^{sticky}, laying off all workers whose MPL falls below the rigid wage. The violet 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; **?**, **?**), which lends external validity to this channel.

[3] This is a robust empirical finding. **?** 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.

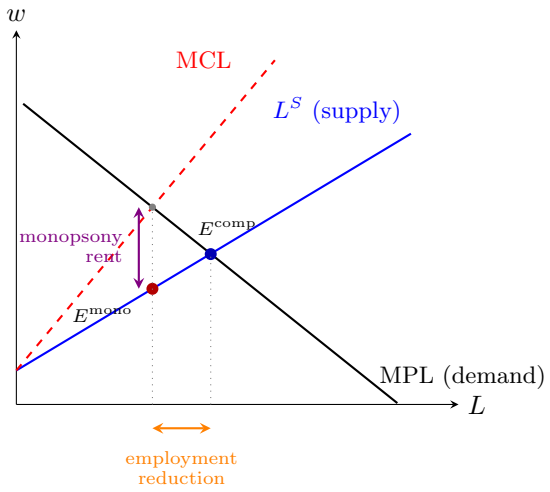

In a competitive market, firms hire until MPL $= w$, reaching E^{comp}. A monopsonist recognizes that hiring an extra worker raises the wage for *all* workers, so the marginal cost of labor (MCL, red dashed) lies above the supply curve. The monopsonist hires where MCL = MPL, reaching E^{mono} with fewer 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. **?** 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 (**?**, **?**) 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 $\text{Var}[(\rfloor \ln I) = \text{Var}[(\rfloor \ln Y)$ 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; **?**'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.

Part III

Problem Sets and Solutions

Subject Index

aggregate risk, 16
Aiyagari model, 175
Alaska Permanent Fund, 147
AR(1) process, 35, 123
Arrow security, 10, 18, 61
Arrow-Debreu equilibrium, 13, 98, 160
autarky, 53

balanced growth path, 87, 89
Banach fixed-point theorem, 101, 102
Bellman equation, 23, 33, 99, 118, 182
Big Push, 109
Blackwell sufficient conditions, 101
borrowing constraint, 19, 30, 176
buffer-stock saving, 163
Bulow-Rogoff theorem, 60
business cycle, 79, 113, 136, 175
Business Cycle Accounting (BCA), 136

cash-in-hand, 37
catch-up growth, 83, 107
Cobb-Douglas production function, 73, 75, 115, 180
coefficient of relative risk aversion, 90, 128, 154
coefficient of variation, 113
competitive equilibrium, 14, 98, 118
complete markets, 10, 19, 31, 46, 70
conditional convergence, 109
contraction mapping, 101
CRRA utility, 147
curse of dimensionality, 49, 103

development accounting, 78
Doob convergence theorem, 35
drifted random walk, 166

efficiency wedge, 136
employment lottery, 131
endogenous grid method (EGM), 187, 188
equity premium puzzle, 177
Euler equation, 32, 96, 127, 137, 149, 184
excess sensitivity, 148

First Welfare Theorem, 15, 32, 98, 118
Frisch elasticity, 113, 136

Great Depression, 143
Great Recession, 131
growth accounting, 78

Hall test, 147
HANK models, 134, 145, 173, 187
Huggett model, 32
human capital, 73, 83, 108

idiosyncratic risk, 17, 29, 60, 187
impulse response function, 122
incomplete markets, 29
indivisible labor, 131
intertemporal elasticity of substitution (IES), 154
intertemporal marginal rate of substitution (IMRS), 22
investment wedge, 136

Kalman filter, 142
Kalman smoother, 142

labor wedge, 134, 136, 138
lack of commitment, 53
Lucas critique, 89

Markov assumption, 166

Markov chain, 23, 34, 49, 117, 121, 178
maximum likelihood estimation, 142
Mincer equation, 77
monopsony, 133

natural debt limit, 10, 29
neoclassical growth model, 90, 92
non-convexities, 106

Pareto optimality, 12, 98
participation constraint, 55
Permanent Income Hypothesis, 147
precautionary saving, 32, 148, 175
pricing kernel, 18
promise-keeping constraint, 53
promised utility, 10

Ramsey-Cass-Koopmans model, 84
Random Walk Hypothesis, 147
randomized controlled trials (RCTs), 111
Real Business Cycle (RBC), 115
recursive competitive equilibrium, 23, 50, 141, 181
Ricardian equivalence, 152
risk sharing, 60

saddle path, 100
Second Welfare Theorem, 16, 118
self-enforcing contract, 55
sequential trading, 10
Solow model, 82, 92, 106, 116
Solow residual, 115
stationary distribution, 29, 69, 179
sticky prices, 96, 119, 137, 188
sticky wages, 132, 138, 188
stochastic discount factor, 10, 118
supermartingale, 34

Tauchen method, 124
total factor productivity (TFP), 73, 93, 108, 115, 138
transversality condition, 92, 161

unconditional convergence, 106

value function iteration (VFI), 101, 106, 119, 183

Author Index

Achdou, Y., 188
Aiyagari, S. R., 175, 187

Barattieri, A., 132
Barro, R. J., 77
Basu, S., 132, 134
Bewley, T. F., 132
Bils, M., 115, 145
Boar, C., 165
Browning, M., 169
Brumberg, R., 147

Campbell, J. Y., 168
Carroll, C. D., 163, 173, 187
Chari, V. V., 134, 142, 143
Christiano, L. J., 144
Cogley, T., 122
Collado, M. D., 169

Davis, J. M., 144

Epstein, L. G., 154

Fernald, J. G., 134
Flavin, M. A., 164
Friedman, M., 147, 148, 155

Gottschalk, P., 132

Hall, R. E., 76, 147, 158, 159, 161, 163, 167
Han, J., 188
Hansen, G. D., 131
Hsieh, C.-T., 170, 172
Huggett, M., 32

Jones, C. I., 76

Karabarbounis, L., 145
Kehoe, P. J., 134, 142, 143
Kehoe, T. J., 144
Keynes, J. M., 148, 155, 156
Kimball, M. S., 134
King, R. G., 115
Klenow, P. J., 115, 145
Krusell, P., 49, 187
Kueng, L., 171, 172
Kydland, F. E., 117

Lasry, J.-M., 188
Lee, J.-W., 77
Lions, P.-L., 188
Lucas, R. E., 90

Malin, B. A., 115, 145
Mankiw, N. G., 168
McGrattan, E. R., 134, 142, 143
Mehra, R., 177
Mincer, J., 77
Modigliani, F., 147
Moll, B., 188
Murphy, K. M., 109

Nason, J. M., 122

Parker, J. A., 169
Patrinos, H. A., 77
Prescott, E. C., 117, 144, 177
Psacharopoulos, G., 77

Rebelo, S. T., 115
Rosenstein-Rodan, P. N., 109

Shleifer, A., 109
Smith, A. A., 49, 187

Souleles, N. S., 169

Tauchen, G., 124

Vishny, R. W., 109

Zeldes, S. P., 164
Zin, S. E., 154