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 3

Final Exam Review Guide

Remark (How to Use This Chapter).

> This chapter is a compressed reference for the final exam. The exam covers two kinds of material in roughly equal proportion:
>
> - **Theory drill** (Section 12.1): Six question types modeled on Problem Set 5 (chapter "Problem Set 8" in this book). Each subsection summarizes the lecture content needed to solve one type of question, with the canonical derivation worked through.
> - **Computation drill** (Section 12.2): Three numerical algorithms from Part II—deterministic Value Function Iteration, stochastic VFI for the RBC model, and the Aiyagari outer loop. The exam asks for the algorithm in writing (steps and rationale), *not* executable code.
>
> The Quick Reference Cheat Sheet in Section 12.3 collects the formulas you should be able to reproduce from memory. The earlier chapters of Part II contain everything in this review and more; this chapter prioritizes only what you need to solve PS5-style problems and write down computational procedures.

3.1 Theory Drill

Each subsection below is organized as **Question** (what an exam-style problem asks), **Setup** (the full model with all primitives stated), **Solution** (the canonical derivation), and **Punchline** (what to remember). The structure mirrors the six problems on Problem Set 5; if you can reproduce these derivations from scratch, you can solve the corresponding exam questions.

3.1.1 Solow Model with Population and TFP Growth

Question. The basic Solow model in the lectures has constant TFP and a constant labor force, and predicts *no* long-run growth. What if we let both TFP and population grow at constant exponential rates? Is there still a balanced growth path (BGP) where capital per worker grows at a constant rate? If so, how does that rate depend on the model's primitives?

Setup. The standard Solow ingredients:

- Cobb–Douglas production: $Y_t = A_t K_t^{\alpha} L_t^{1-\alpha}$, with $\alpha \in (0,1)$.
- Constant saving rate $s \in (0,1)$: investment $I_t = sY_t$.
- Capital accumulation: $K_{t+1} = I_t + (1-\delta)K_t$, with $\delta \in (0,1)$.
- Two new ingredients: TFP grows at rate g_A, i.e., $A_t = A_0 e^{g_A t}$; labor grows at rate g_L, i.e., $L_t = L_0 e^{g_L t}$.

You are asked to (i) show that a BGP for capital exists, and (ii) compute the BGP growth rates of total capital g_K and capital per worker g_k.

Solution. The technique is **guess and verify**: conjecture that on the BGP, $k_t \equiv K_t/L_t$ grows at some constant rate g_k, then plug the conjecture into the law of motion and demand consistency.

Step 1: Per-worker law of motion. Divide $K_{t+1} = sA_t K_t^{\alpha} L_t^{1-\alpha} + (1-\delta)K_t$ by L_{t+1} and use $L_{t+1}/L_t = e^{g_L}$:

$$e^{g_L} k_{t+1} \; = \; sA_t\, k_t^{\alpha} + (1-\delta)\, k_t. \tag{$*$}$$

Step 2: Substitute the BGP guess. Plug $k_t = k_0 e^{g_k t}$ and $A_t = A_0 e^{g_A t}$ into ($*$), divide both sides by $k_0 e^{g_k t}$:

$$e^{g_L + g_k} \; = \; sA_0\, k_0^{\alpha-1}\, e^{(g_A + (\alpha-1)g_k)\, t} + (1-\delta).$$

Step 3: Demand time-invariance. The left side is constant in t. For the right side to also be constant, the coefficient of t in the exponent must vanish:

$$g_A + (\alpha-1)\, g_k = 0 \implies \boxed{g_k = \frac{g_A}{1-\alpha}, \qquad g_K = g_L + g_k = g_L + \frac{g_A}{1-\alpha}.}$$

Punchline.

- TFP growth is amplified into per-worker capital growth by the multiplier $1/(1-\alpha) > 1$. The mechanism: a TFP gain raises the marginal product of capital, calling forth additional investment, which raises the marginal product further—a within-period feedback that the Cobb–Douglas exponent α controls.
- Population growth contributes one-for-one to *total* capital ($g_K = g_L + \cdots$) but not at all to capital *per worker*. "Spreading more capital across more workers" offsets the additional accumulation, leaving k_t unaffected.
- The technique (substitute exponential guess, force the t-exponent to zero) is the standard tool for any BGP question. Memorize it.

3.1.2 Two-Period Saving with Risky Returns

Question. A consumer lives for two periods with log utility. They earn labor income in either period 1 or period 2 (two distinct sub-questions), and the gross return $1+r$ they face on saving/borrowing may be either certain or random. The exam asks: does the risk in r

matter for first-period consumption C_1? In particular, if r becomes uncertain while keeping $\mathbb{E}(r)$ fixed (a mean-preserving spread), does C_1 rise, fall, or stay the same?

The economic content: this is a question about precautionary saving versus the income/substitution-effect cancellation. Whether C_1 moves depends on a delicate interplay between (i) when income arrives and (ii) the curvature of the utility function in the random variable r.

Setup. Lifetime utility $U = \ln C_1 + \mathbb{E}[\ln C_2]$. Two scenarios:

Case A: income upfront. Endowment $Y_1 > 0$ in period 1 and zero in period 2. The agent saves $Y_1 - C_1$ at gross return $1 + r$, so $C_2 = (1 + r)(Y_1 - C_1)$.

Case B: income back-loaded. Endowment 0 in period 1 and $Y_2 > 0$ in period 2. The agent borrows C_1 in period 1 and repays $(1+r)C_1$ in period 2, so $C_2 = Y_2 - (1+r)C_1$.

In both cases, derive (i) the FOC for C_1, and (ii) the response of C_1 to a mean-preserving spread in r.

Solution: Case A (saver).

Lifetime utility:

$$U \;=\; \ln C_1 + \mathbb{E}[\ln(1 + r)(Y_1 - C_1)] \;=\; \ln C_1 + \ln(Y_1 - C_1) + \mathbb{E}[\ln(1 + r)].$$

The crucial observation: log utility makes the random variable r enter *additively* through $\mathbb{E}[\ln(1 + r)]$, which does not depend on C_1. The FOC therefore drops r entirely:

$$\frac{1}{C_1} \;=\; \frac{1}{Y_1 - C_1} \;\Longrightarrow\; C_1 = \tfrac{1}{2}Y_1.$$

A mean-preserving spread in r changes only $\mathbb{E}[\ln(1+r)]$, which the FOC ignores. Hence C_1 is *unchanged.*

Solution: Case B (borrower).

Lifetime utility:

$$U \;=\; \ln C_1 + \mathbb{E}\big[\ln(Y_2 - (1 + r)C_1)\big].$$

Now r multiplies C_1 inside $\ln(\cdot)$, so r enters the FOC:

$$\frac{1}{C_1} \;=\; \mathbb{E}\left[\frac{1 + r}{Y_2 - (1 + r)C_1}\right] \;\equiv\; \mathbb{E}[f(r)].$$

Compute the curvature of f. By the quotient rule:

$$f'(r) = \frac{Y_2}{[Y_2 - (1 + r)C_1]^2}, \qquad f''(r) = \frac{2Y_2C_1}{[Y_2 - (1 + r)C_1]^3} > 0,$$

provided $C_2 > 0$ in every state. By Jensen's inequality, an MPS in r at constant $\mathbb{E}(r)$ strictly raises $\mathbb{E}[f(r)]$. The FOC then forces $1/C_1$ to rise—i.e., C_1 *falls.*

Punchline.

- In **Case A**, log utility and saving combine to make C_1 independent of return uncertainty. The substitution effect (higher expected return ⇒ save more) and income effect (higher lifetime wealth ⇒ consume more) cancel exactly under log preferences. *No precautionary response.*

- In **Case B**, the agent is borrowing, and uncertainty in r makes the marginal cost of borrowing $f(r) = (1 + r)/(Y_2 - (1 + r)C_1)$ stochastic. Convexity of f in r means the bad states (high r, very low C_2, very high marginal utility) outweigh the good states. *Borrowing falls; C_1 strictly decreases.*

- **General lesson for the exam.** The direction of the response to MPS in r depends on whether r enters the FOC additively (Case A; no effect on C_1) or multiplicatively (Case B; effect via Jensen on a convex function).

3.1.3 RBC: Bellman, Euler, and Linear-Quadratic Guess-and-Verify

Question. This subsection answers two questions a typical RBC problem combines.

(1) **General theory.** Given a stochastic neoclassical growth model with state (a, k), write down the Bellman equation, derive the Euler equation, and (under quadratic utility) show that consumption follows Hall's random-walk hypothesis.

(2) **Closed-form solution.** For a particular linear-quadratic specification (Blanchard–Fischer), guess that consumption is linear in capital and the productivity shock, and solve for the coefficients of the policy function.

Setup for (1): General stochastic Bellman. A planner faces:

- Per-period utility $u(c)$ (smooth, concave, satisfying Inada conditions).
- Cobb–Douglas production with stochastic TFP: $y_t = e^{a_t} f(k_t)$ where a_t follows an AR(1).
- Standard capital accumulation: $k_{t+1} = e^{a_t} f(k_t) + (1 - \delta)k_t - c_t$.
- Discount factor $\beta \in (0, 1)$.

Solution for (1): Bellman, FOC, Euler.

$$V(a, k) \;=\; \max_{c,k'} \big\{u(c) + \beta\,\mathbb{E}[V(a', k') \mid a]\big\}, \quad k' = e^a f(k) + (1 - \delta)k - c.$$

FOC w.r.t. k'. The first-order condition (using $\partial k'/\partial c = -1$):

$$u'(c) \;=\; \beta\,\mathbb{E}[V_{k'}(a', k') \mid a].$$

Envelope theorem. Differentiating the Bellman in k:

$$V_k(a, k) \;=\; u'(c) \cdot \big[e^a f'(k) + 1 - \delta\big].$$

Combining. Forward-shift the envelope ($V_{k'}(a', k') = u'(c')\cdot[e^{a'} f'(k')+1-\delta]$) and substitute into the FOC:

$$\boxed{u'(c_t) \;=\; \beta\,\mathbb{E}_t\big[u'(c_{t+1})\,(e^{a_{t+1}} f'(k_{t+1}) + 1 - \delta)\big].}$$

Hall's RWH (special case). If $u(c) = c - \theta c^2$ (so $u'(c) = 1 - 2\theta c$ is linear in c) and the gross return is constant at $1 + r$ with $\beta(1 + r) = 1$, the Euler equation collapses:

$$1 - 2\theta c_t \;=\; \mathbb{E}_t[1 - 2\theta c_{t+1}] \iff c_t = \mathbb{E}_t[c_{t+1}].$$

Consumption is a martingale; changes from c_t to c_{t+1} are unforecastable from time-t information.

Setup for (2): Linear-quadratic RBC (Blanchard–Fischer). A specific specification where the Euler simplifies:

- Period utility $u(C) = C - \theta C^2$, $\theta > 0$.
- Discount: the household discounts at rate $\rho > 0$, so the discount factor is $1/(1+\rho)$.
- Linear production with additive shock: $Y_t = AK_t + e_t$, with $A = \rho$ (a parameter restriction that simplifies the algebra).
- No depreciation: $K_{t+1} = K_t + Y_t - C_t = (1+\rho)K_t + e_t - C_t$.
- AR(1) shock: $e_t = \phi e_{t-1} + \varepsilon_t$, with $|\phi| < 1$ and ε_t i.i.d. mean-zero.

Question: guess $C_t = \alpha + \beta K_t + \gamma e_t$, and solve for (α, β, γ).

Solution for (2): Guess and verify. The Euler equation in this LQ environment simplifies to $C_t = \mathbb{E}_t[C_{t+1}]$ (Hall's RWH from above, with $\beta(1+r) = 1$ effectively imposed by $A = \rho$).

Step 1: Substitute the linear policy guess into the resource constraint. If $C_t = \alpha + \beta K_t + \gamma e_t$,

$$K_{t+1} \;=\; (1+\rho)K_t + e_t - (\alpha + \beta K_t + \gamma e_t) \;=\; (1+\rho-\beta)K_t + (1-\gamma)e_t - \alpha.$$

Step 2: Substitute into the Euler $C_t = \mathbb{E}_t[C_{t+1}]$.

$$\alpha + \beta K_t + \gamma e_t \;=\; \mathbb{E}_t[\alpha + \beta K_{t+1} + \gamma e_{t+1}] \;=\; \alpha + \beta\,\mathbb{E}_t[K_{t+1}] + \gamma\phi e_t.$$

Plug in $\mathbb{E}_t[K_{t+1}] = (1+\rho-\beta)K_t + (1-\gamma)e_t - \alpha$:

$$\alpha + \beta K_t + \gamma e_t \;=\; \alpha + \beta\big[(1+\rho-\beta)K_t + (1-\gamma)e_t - \alpha\big] + \gamma\phi e_t.$$

Step 3: Match coefficients.

- *K_t:* $\beta = \beta(1+\rho-\beta) \Longrightarrow \beta = \rho$ (the non-trivial root).
- *e_t:* $\gamma = \beta(1-\gamma) + \gamma\phi \Longrightarrow \gamma(1+\beta-\phi) = \beta \Longrightarrow \gamma = \rho/(1+\rho-\phi)$.
- *Constant:* $\alpha = \alpha - \beta\alpha \Longrightarrow \alpha = 0$ (since $\beta = \rho > 0$).

Result: $C_t = \rho K_t + \frac{\rho}{1+\rho-\phi}\, e_t$.

Punchline.

- The general FOC and envelope theorem deliver the Euler equation in any neoclassical model. Memorize the procedure: FOC w.r.t. k', then envelope theorem, then forward-shift.
- Hall's random-walk result requires both quadratic utility and $\beta(1+r) = 1$. Quadratic alone makes u' linear; $\beta(1+r) = 1$ kills the $\beta(1+r)$ factor in the Euler equation.
- For LQ-RBC, the policy function is linear in the state. The four-step procedure (guess → resource → Euler → match coefficients) is exam-mechanical—no creativity required, only careful algebra.

3.1.4 PIH and the Farmer/Non-Farmer Empirical Puzzle

Question. The empirical observation: in cross-sectional household data, farmers' average income is lower than non-farmers', and farmer income fluctuates more from year to year. If you ran a regression of consumption C_i on current income Y_i separately for the two groups, the regression slopes would differ. The exam asks: under the permanent-income hypothesis (PIH), what does theory predict about how the two slopes differ, and how do we explain the prediction?

The framing matters for the exam: the question is about *the apparent MPC*, the slope $\hat{b}$ in a regression of C_i on Y_i. PIH does *not* say that the apparent MPC is constant across groups—in fact, it predicts the apparent MPC depends on the income process.

Setup. Decompose individual income into permanent and transitory components:

$$Y_i \;=\; Y_i^P + Y_i^T, \qquad \mathrm{Cov}[(,Y]_i^P, Y_i^T) = 0.$$

Under the PIH, consumption is a constant fraction of permanent income:

$$C_i \;=\; \phi\, Y_i^P, \qquad \phi = r/(1+r) \text{ in the infinite-horizon case.}$$

You are asked to compute the cross-sectional regression slope $\hat{b} = \mathrm{Cov}[(,Y]_i, C_i)/\,\mathrm{Var}[(]\,Y_i)$ and interpret it for farmers vs. non-farmers.

Solution. Direct computation:

$$\hat{b} \;=\; \frac{\mathrm{Cov}[(,Y]_i, C_i)}{\mathrm{Var}[(]\,Y_i)} \;=\; \frac{\mathrm{Cov}[(,Y]_i^P + Y_i^T,\, \phi Y_i^P)}{\mathrm{Var}[(]\,Y_i^P) + \mathrm{Var}[(]\,Y_i^T)} \;=\; \phi \cdot \frac{\mathrm{Var}[(]\,Y_i^P)}{\mathrm{Var}[(]\,Y_i^P) + \mathrm{Var}[(]\,Y_i^T)},$$

using $\mathrm{Cov}[(,Y]^P, Y^T) = 0$. Up to the constant factor ϕ (close to 1 in the infinite-horizon limit), the regression slope equals *the share of permanent variance in total income variance.*

Application.

- **Farmers.** Their incomes are dominated by transitory shocks (weather, crop prices), so $\mathrm{Var}[(]\,Y^T)$ is large relative to $\mathrm{Var}[(]\,Y^P)$. The signal-to-total ratio is small, hence $\hat{b}$ is small. The estimated consumption function looks flat, and the "apparent MPC" is low.
- **Non-farmers.** Wage income is much more stable, with $\mathrm{Var}[(]\,Y^T)$ small relative to $\mathrm{Var}[(]\,Y^P)$. The signal-to-total ratio is close to 1, so $\hat{b} \approx \phi \approx 1$. The estimated consumption function looks steep, with apparent MPC close to 1.

Punchline.

- A naive observer comparing the two slopes might conclude that farmers "don't respond" to income while non-farmers respond strongly. The PIH explanation is exactly the opposite: *both* groups respond identically to permanent income changes; the regression slope is small for farmers because their measured income is mostly noise that they correctly ignore.
- The key formula $\hat{b} = \phi \cdot \mathrm{Var}[(]\,Y^P)/[\mathrm{Var}[(]\,Y^P) + \mathrm{Var}[(]\,Y^T)]$ is exam-standard for any cross-section question about apparent MPC vs. structural MPC.

3.1.5 CRRA Consumption with Lognormal Shocks

Question. A consumer has CRRA utility and faces stochastic future consumption. The exam asks for four things in sequence:

(a) Write down the Euler equation.

(b) Assume conditional lognormality of C_{t+1} and rewrite the Euler equation in terms of $\mathbb{E}_t[\ln C_{t+1}]$, $\ln C_t$, σ^2, and primitives.

(c) Show that this implies $\ln C$ follows a random walk with drift.

(d) Interpret how the drift depends on σ^2, in light of the precautionary saving motive.

The economic point is to show that under CRRA (which has $u''' > 0$), uncertainty about future consumption raises expected consumption growth—a smoking-gun signature of precautionary saving that distinguishes CRRA from quadratic utility.

Setup.

- Utility $u(C) = C^{1-\theta}/(1-\theta)$, $\theta > 0$, $\theta \neq 1$. (Marginal utility $u'(C) = C^{-\theta}$.)
- Constant interest rate r, but $\beta(1+r) \neq 1$ in general.
- The conditional distribution of $\ln C_{t+1}$ given time-t information is normal with conditional mean $\mu_t = \mathbb{E}_t[\ln C_{t+1}]$ and conditional variance σ^2.

Solution.

Step 1: CRRA Euler equation. From $u'(C_t) = \beta(1+r)\mathbb{E}_t[u'(C_{t+1})]$:

$$C_t^{-\theta} \;=\; \beta(1+r)\,\mathbb{E}_t\big[C_{t+1}^{-\theta}\big].$$

Step 2: Lognormal substitution. If $\ln C_{t+1} \sim \mathcal{N}(\mu_t, \sigma^2)$ conditional on time-t, then $-\theta \ln C_{t+1} \sim \mathcal{N}(-\theta\mu_t, \theta^2\sigma^2)$, and the lognormal MGF identity $\mathbb{E}[e^x] = e^{\mu+V/2}$ (for $x \sim \mathcal{N}(\mu, V)$) gives

$$\mathbb{E}_t\big[C_{t+1}^{-\theta}\big] \;=\; \mathbb{E}_t\big[e^{-\theta \ln C_{t+1}}\big] \;=\; \exp\Big(-\theta\mu_t + \tfrac{\theta^2\sigma^2}{2}\Big).$$

Substitute into the Euler equation and take logs:

$$-\theta \ln C_t \;=\; \ln\beta + \ln(1+r) - \theta\mu_t + \tfrac{\theta^2\sigma^2}{2}.$$

Step 3: Drifted random walk. Solve for μ_t:

$$\boxed{\mathbb{E}_t[\ln C_{t+1}] \;=\; \ln C_t + \underbrace{\frac{\ln\beta + \ln(1+r)}{\theta}}_{\text{intertemporal substitution}} + \underbrace{\frac{\theta\sigma^2}{2}}_{\text{precautionary}}.}$$

Defining the innovation $u_{t+1} \equiv \ln C_{t+1} - \mathbb{E}_t[\ln C_{t+1}]$ (which has $\mathbb{E}_t[u_{t+1}] = 0$), we have

$$\ln C_{t+1} \;=\; \ln C_t + a + u_{t+1}, \qquad a \equiv \tfrac{\ln\beta(1+r)}{\theta} + \tfrac{\theta\sigma^2}{2},$$

i.e., $\ln C_t$ is a random walk with drift a.

Step 4: Interpretation. The drift has two pieces:

- **Intertemporal-substitution channel** $\frac{\ln \beta(1+r)}{\theta}$: positive when $\beta(1+r) > 1$ (the market rewards saving more than the household discounts the future), negative otherwise. With $\beta(1+r) = 1$, this term vanishes.
- **Precautionary channel** $\frac{\theta\sigma^2}{2}$: *always positive when* $\theta > 0$. A higher σ^2 raises expected consumption growth—the household saves more today (depressing C_t), which mechanically lifts $\mathbb{E}_t[\Delta \ln C_{t+1}]$.

Punchline.

- The lognormal trick ($\mathbb{E}[e^x] = e^{\mu+V/2}$) converts the multiplicative Euler equation into an *additive* relationship in log consumption, exposing the random-walk-with-drift structure.
- The precautionary term $\theta\sigma^2/2$ exists precisely because $u''' > 0$ for CRRA. Under quadratic utility ($u''' = 0$), this term would vanish, and consumption growth would not respond to uncertainty—a smoking gun for distinguishing the two utility classes.
- Memorize the lognormal MGF identity. It is the single technical tool that makes this problem solvable in closed form.

3.1.6 Random-Walk Consumption: Martingale Property and MPC Out of AR(1) Shocks

Question. A consumer follows the Hall-style consumption rule

$$C_t \;=\; \frac{r}{1+r}\big[A_t + H_t\big], \qquad H_t \equiv \sum_{s=0}^{\infty} \frac{\mathbb{E}_t[Y_{t+s}]}{(1+r)^s},$$

where A_t is non-state-contingent assets and H_t is human wealth (the present value of expected income). The exam asks two parts:

(a) Show that $\mathbb{E}_t[C_{t+1}] = C_t$ (consumption is a martingale).

(b) Suppose income follows an AR(1) process $Y_t = \phi Y_{t-1} + u_t$ with i.i.d. innovations u_t. If a unit shock $u_t = 1$ arrives, by how much does consumption C_t change?

The economic content: (a) tests your ability to derive the random-walk property using the law of iterated expectations on human wealth; (b) tests whether you can compute the marginal propensity to consume out of a persistent income shock.

Setup.

- The consumption rule above (you may take this as given; it is derived from the PIH theorem).
- Asset evolution: $A_{t+1} = (1+r)(A_t + Y_t - C_t)$.
- Income process for part (b): $Y_t = \phi Y_{t-1} + u_t$ with $|\phi| < 1$ and $\mathbb{E}_{t-1}[u_t] = 0$.

Solution to (a): Consumption is a martingale.

The key step is computing $\mathbb{E}_t[H_{t+1}]$. By the law of iterated expectations, $\mathbb{E}_t[\mathbb{E}_{t+1}[Y_{t+1+s}]] = \mathbb{E}_t[Y_{t+1+s}]$, so

$$\mathbb{E}_t[H_{t+1}] \;=\; \mathbb{E}_t\left[\sum_{s=0}^{\infty} \frac{\mathbb{E}_{t+1}[Y_{t+1+s}]}{(1+r)^s}\right] \;=\; \sum_{s=0}^{\infty} \frac{\mathbb{E}_t[Y_{t+1+s}]}{(1+r)^s}.$$

Reindex with $s' = s+1$ to align with the definition of H_t:

$$\mathbb{E}_t[H_{t+1}] \;=\; (1+r)\sum_{s'=1}^{\infty} \frac{\mathbb{E}_t[Y_{t+s'}]}{(1+r)^{s'}} \;=\; (1+r)\big(H_t - Y_t\big).$$

(The last equality is H_t minus its $s=0$ term.) This is the **key lemma**.

Combine with the asset law $\mathbb{E}_t[A_{t+1}] = (1+r)(A_t + Y_t - C_t)$:

$$\mathbb{E}_t[A_{t+1} + H_{t+1}] \;=\; (1+r)(A_t + Y_t - C_t) + (1+r)(H_t - Y_t) \;=\; (1+r)(A_t + H_t - C_t).$$

Therefore

$$\mathbb{E}_t[C_{t+1}] \;=\; \frac{r}{1+r}\,(1+r)(A_t + H_t - C_t) \;=\; r(A_t + H_t) - rC_t.$$

Substitute $r(A_t + H_t) = (1+r)C_t$ from the consumption rule:

$$\mathbb{E}_t[C_{t+1}] \;=\; (1+r)C_t - rC_t \;=\; C_t. \qquad \square$$

Solution to (b): MPC out of an AR(1) innovation.

A unit innovation $u_t = 1$ at time t raises $\mathbb{E}_t[Y_{t+s}]$ by ϕ^s for each $s \geq 0$ (start from $\phi^0 = 1$ at $s = 0$, then ϕ^1 at $s = 1$, etc.). The change in human wealth is a geometric series:

$$\Delta H_t \;=\; \sum_{s=0}^{\infty} \frac{\phi^s}{(1+r)^s} \;=\; \frac{1}{1 - \phi/(1+r)} \;=\; \frac{1+r}{1+r-\phi}.$$

The change in consumption follows from $\Delta C_t = \frac{r}{1+r}\Delta H_t$ (assets A_t are unchanged at the moment of the shock):

$$\boxed{\Delta C_t \;=\; \frac{r}{1+r}\cdot\frac{1+r}{1+r-\phi} \;=\; \frac{r}{1+r-\phi}.}$$

Limit cases.

- $\phi \to 0$ (i.i.d. income): MPC $\to r/(1+r) \approx r$ for small r. The agent treats a one-period windfall as adding only its annuity value to lifetime resources.
- $\phi \to 1$ (permanent shock): MPC $\to 1$. A permanent change in income raises consumption nearly one-for-one—the textbook PIH response.
- Intermediate ϕ: the MPC interpolates between the two limits, reflecting how persistent the shock is expected to be.

Punchline.

- **Random-walk derivation.** The trick is to first establish the lemma $\mathbb{E}_t[H_{t+1}] = (1 + r)(H_t - Y_t)$, then combine with the asset law. Don't try to manipulate $\mathbb{E}_t[C_{t+1}]$ directly without this lemma—it gets messy.

- **MPC formula.** $\text{MPC} = r/(1 + r - \phi)$ is monotone increasing in ϕ. It interpolates smoothly between $\sim r$ (i.i.d.) and 1 (permanent), reflecting how much of an income shock is expected to persist into the future.

3.2 Computation Drill

The exam will ask you to write down algorithm steps, not actual code. Each algorithm below should be reproduced from memory: the data structures (grids), the update rule (Bellman or pushforward), the convergence criterion, and—when relevant—the theoretical guarantee (contraction mapping for VFI; fixed point for stationary distribution; market clearing for Aiyagari).

3.2.1 Value Function Iteration (Deterministic Neoclassical Model)

Algorithm: Deterministic VFI

Solving $V(k) = \max_{k'}\{u(Af(k) + (1-\delta)k - k') + \beta V(k')\}$ on a discretized state space.

Step 1: **Discretize the state space.** Choose a grid $\mathcal{K} = \{k_1, k_2, \ldots, k_N\}$ for capital. Common choice: $N = 500$ points uniformly on $[k_{\min}, k_{\max}]$ where $[k_{\min}, k_{\max}]$ brackets the steady state k^*. Use a log-spaced grid if precision is needed near $k = 0$ (Inada singularity).

Step 2: **Initial guess.** $V^{(0)}(k) = 0$ for all $k \in \mathcal{K}$. Any bounded initial guess converges to the same fixed point by the Banach Fixed Point Theorem.

Step 3: **Bellman update.** For each $k \in \mathcal{K}$, compute

$$V^{(i+1)}(k) \;=\; \max_{k' \in \mathcal{K}\,:\,c>0} \Big\{ u\big(Af(k) + (1-\delta)k - k'\big) + \beta V^{(i)}(k') \Big\}.$$

The constraint $c > 0$ requires $k' < Af(k) + (1-\delta)k$; values violating it are excluded from the maximand. Also record the optimizer at each grid point.

Step 4: **Convergence check.** Stop when $\max_k |V^{(i+1)}(k) - V^{(i)}(k)| < \varepsilon$ (typically $\varepsilon = 10^{-6}$). Otherwise increment i and return to Step 3. Always set a maximum iteration count (e.g., 1000) to guard against runaway loops.

Step 5: **Extract policy.** After convergence, the optimal policy is $g(k) = \arg\max_{k'}$ from the final iteration. Apply iteratively to track transition paths from any initial k_0.

Why it works. The Bellman operator T defined by $(TV)(k) = \max_{k'}\{u(\cdot) + \beta V(k')\}$ is a contraction mapping with modulus $\beta < 1$ on the space of bounded continuous functions.

By the Banach Fixed Point Theorem there exists a unique fixed point V^*, and $V^{(i)} \to V^*$ uniformly from *any* bounded initial guess. This is what justifies starting at $V^{(0)} = 0$.

Common pitfalls.

- *Log/CRRA singularity at $c = 0$.* Always enforce $k' < Af(k) + (1-\delta)k$ strictly to avoid $\ln(0) = -\infty$ or $C^{-\theta} \to \infty$.
- *Grid endpoint binding.* If the optimal k' is consistently at the maximum grid value, the grid is too narrow; widen it.
- *Convergence diagnostics.* Plot the maximum change $\max_k |V^{(i+1)} - V^{(i)}|$ across iterations; it should fall geometrically at rate β.

3.2.2 Stochastic VFI for the Real Business Cycle Model

Algorithm: Stochastic VFI for RBC

Solving $V(a,k) = \max_{k'}\{u(c) + \beta\,\mathbb{E}[V(a',k') \mid a]\}$ where a follows an AR(1).

Step 1: Discretize the AR(1) process. If $a' = \rho a + \varepsilon$, $\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon^2)$, use the **Tauchen method**: choose M grid points $\{a^1, \dots, a^M\}$ spanning ±3 unconditional standard deviations of a. Compute the $M \times M$ transition matrix Π with entries

$$\pi_{ij} \;=\; \Phi\!\left(\frac{a^j + \Delta/2 - \rho a^i}{\sigma_\varepsilon}\right) - \Phi\!\left(\frac{a^j - \Delta/2 - \rho a^i}{\sigma_\varepsilon}\right),$$

where Δ is the grid spacing and Φ is the standard normal CDF. As $M \to \infty$, the discrete chain converges to the AR(1).

Step 2: Discretize capital. Choose $\mathcal{K} = \{k_1, \dots, k_N\}$ as in the deterministic case.

Step 3: Bellman update. For each $(a,k) \in \{a^1, \dots, a^M\} \times \mathcal{K}$:

$$V^{(i+1)}(a,k) = \max_{k' \in \mathcal{K}\,:\,c>0} \left\{ u\big(e^a f(k) + (1-\delta)k - k'\big) + \beta \sum_{a'} \pi(a' \mid a)\, V^{(i)}(a',k') \right\}.$$

Record the optimizer.

Step 4: Convergence check. Stop when $\max_{a,k} |V^{(i+1)} - V^{(i)}| < \varepsilon$.

Step 5: Simulate. Starting from (a_0, k_0), draw a path $\{a_t\}_{t=1}^T$ from the Markov chain, then iterate $k_{t+1} = g(a_t, k_t)$ using the converged policy. Compute realized $\{y_t = e^{a_t} f(k_t),\, c_t,\, i_t = k_{t+1} - (1-\delta)k_t\}$.

Step 6: Compute moments. Report coefficients of variation: $\mathrm{CV}(x) = \sqrt{\mathrm{Var}[(]\,x)}/\mathbb{E}(x)$, comparing $\mathrm{CV}(c) < \mathrm{CV}(y) < \mathrm{CV}(i)$ to confirm the model reproduces the canonical investment-volatility fact.

Why CV rather than raw variance. The series $\{c, y, i\}$ have very different mean levels, so raw variance comparisons confound scale with volatility. The coefficient of variation

$\mathrm{CV}(x) = \sqrt{\mathrm{Var}[(]\,x)}/\mathbb{E}(x)$ is unit-free and is the standard object reported in the RBC literature.

3.2.3 Aiyagari Outer Loop (Three Nested Loops)

Algorithm: Aiyagari Stationary Equilibrium

Find r^* such that the household sector's aggregate savings equal the firm's capital demand.

Step 1: Outer loop initial guess. Set r_0 slightly below $1/\beta - 1$ (since precautionary motives push r^* below the representative-agent benchmark).

Step 2: Compute w and K^{firm} from firm FOCs. For Cobb-Douglas $F(K, L) = K^\alpha L^{1-\alpha}$:

$$\frac{K}{L} = \left(\frac{\alpha}{r_0}\right)^{1/(1-\alpha)}, \qquad w = (1-\alpha)(K/L)^\alpha, \qquad K^{\text{firm}} = L \cdot (K/L).$$

Here $L = \mathbb{E}[\exp(z)]$ is the constant aggregate labor supply.

Step 3: Solve household problem (middle loop). Given (r_0, w), run VFI on the discretized (z, k) state space to obtain the policy function $g_k(z, k)$. The Bellman is

$$V(z,k) = \max_{k'} \left\{u(c) + \beta\, \mathbb{E}[V(z',k') \mid z]\right\}, \quad c = (1+r_0-\delta)k + we^z - k', \quad k' \geq -B.$$

Step 4: Solve stationary distribution (inner loop). Given g_k, find the fixed point Λ^* of the pushforward operator:

- Initialize $\Lambda^{(0)}$ uniform on the grid.
- Iterate $\Lambda^{(j+1)}(z', k') = \sum_{z,k} \Lambda^{(j)}(z,k) \cdot \pi(z' \mid z) \cdot \mathbf{1}\{g_k(z,k) = k'\}$.
- Off-grid lottery: when $g_k(z,k)$ falls between grid points $k_{j'}$ and $k_{j'+1}$, split the mass linearly: send fraction $1-\theta$ to $k_{j'}$ and θ to $k_{j'+1}$, where θ is the relative position. This preserves total mass exactly.
- Stop when $\|\Lambda^{(j+1)} - \Lambda^{(j)}\| < \varepsilon$.

Step 5: Compute excess demand.

$$X(r_0) \;=\; K^{\text{firm}}(r_0) \;-\; \int g_k(z,k)\, d\Lambda^*(z,k) \;=\; K^{\text{firm}} - K^{\text{HH}}.$$

Step 6: Convergence check on outer loop.

- If $|X(r_0)| < \varepsilon_r$, stop. Equilibrium is $r^* = r_0$.
- If $|X(r_0)| \geq \varepsilon_r$, update via bisection (if monotonicity is known) or relaxation $r_0 \leftarrow r_0 - \kappa \cdot X(r_0)$ for small $\kappa > 0$. Return to Step 2.

Economic interpretation. The outer loop traces out the intersection of the household sector's upward-sloping capital supply $K^{\text{HH}}(r)$ (precautionary saving rises with r) and the firm sector's downward-sloping capital demand $K^{\text{firm}}(r) = L(\alpha/r)^{1/(1-\alpha)}$. At the equilibrium r^*, the two intersect.

The atom at the borrowing constraint. A robust outcome of the algorithm is that the stationary distribution Λ^* has an atom at $k = -B$: the mass of households whose policy $g_k(z, k)$ would set $k' < -B$ in the absence of the constraint and so are forced to $k' = -B$ exactly. This atom is the source of wealth inequality at the bottom of the distribution, distinct from the long right tail generated by persistently lucky high-z households.

3.3 Quick Reference Cheat Sheet

Remark (Formulas to Reproduce From Memory).

Result	**Formula**
Solow BGP per-worker growth	$g_k = g_A/(1-\alpha)$
Solow BGP total-capital growth	$g_K = g_L + g_A/(1-\alpha)$
Neoclassical Euler	$u'(C_t) = \beta\, u'(C_{t+1})\,[F_K(K_{t+1}, L) + 1 - \delta]$
Hall's RWH (quadratic u, $\beta(1+r) = 1$)	$C_t = \mathbb{E}_t[C_{t+1}]$
CRRA + lognormal drift	$\mathbb{E}_t[\ln C_{t+1}] - \ln C_t = \frac{\ln \beta(1+r)}{\theta} + \frac{\theta\sigma^2}{2}$
PIH consumption rule	$C = \phi \cdot \text{PI}, \quad \phi = r/(1+r)$ (infinite horizon)
PIH MPC out of AR(1) shock	$\text{MPC} = r/(1 + r - \phi)$
PIH cross-section regression slope	$\hat{b} = \text{Var}(Y^P)/[\text{Var}(Y^P) + \text{Var}(Y^T)]$
RBC (LQ) policy function	$C_t = \rho K_t + \frac{\rho}{1+\rho-\phi} e_t$
Banach contraction guarantee	$V^{(i+1)} = TV^{(i)}$ converges if T is contraction with modulus < 1

Remark (Common Exam Tricks).

- **Log utility cancels income/sub effects.** When labor income is upfront and the agent saves at random return r, C_1 is independent of the distribution of r—the substitution effect (higher $r \Rightarrow$ save more) and income effect (higher lifetime wealth $\Rightarrow$ consume more) exactly cancel. Recognize this whenever you see log utility plus saving.
- **Jensen's inequality for mean-preserving spreads.** If a function f is convex (resp. concave), an MPS in the random variable raises (resp. lowers) $\mathbb{E}[f]$. To apply this to a comparative-static question, identify the relevant f inside the FOC and check the sign of f'' by direct differentiation.
- **Guess-and-verify for linear-quadratic problems.** The closed-form solution to LQ

Bellman problems is always linear in the state. Guess $C_t = \alpha + \beta K_t + \gamma e_t$, substitute into the resource constraint to express K_{t+1} in terms of K_t, e_t, then substitute into the Euler equation and match coefficients. The constant (α), capital-loading (β), and shock-loading (γ) are pinned down independently.

- **Lognormal MGF.** If $x \sim \mathcal{N}(\mu, V)$, then $\mathbb{E}[e^x] = e^{\mu+V/2}$. Use this whenever the Euler equation contains $\mathbb{E}[C^{-\theta}]$ and $\ln C$ is conditionally normal.
- **Random-walk $\Rightarrow$ AR(1) MPC.** The PIH consumption response to an AR(1) income innovation with persistence ϕ is $r/(1 + r - \phi)$. Limits: $\phi = 0 \Rightarrow$ MPC $\approx r$; $\phi = 1 \Rightarrow$ MPC $\to 1$.
- **Computational algorithms have a common skeleton.** Discretize state, guess initial value, update with Bellman or pushforward, check convergence, extract policy / aggregate. Variations: deterministic vs. stochastic (add a discrete-shock dimension and an expectation in the Bellman); Aiyagari adds an outer loop on r.

Part III

Problem Sets and Solutions

Subject Index

Author Index