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; $\pi(s^\tau \vert s^t)$ is conditional.
ε_t	Innovation / shock realization.
ρ	Persistence parameter of an AR(1) process; $\rho = \psi$ in Ch. 2's CARA example.
Endowment and production	
$y(s^t)$, Y_t	Stochastic endowment; aggregate output.

(continued on next page)

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

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

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

Part I

Heterogeneous Agents in Complete and Incomplete Markets

Lectures by Maria-Jose Carreras-Valle

Chapter 1

Complete Markets under Uncertainty

Remark (Notation in This Chapter).

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

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

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

Assume:

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

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

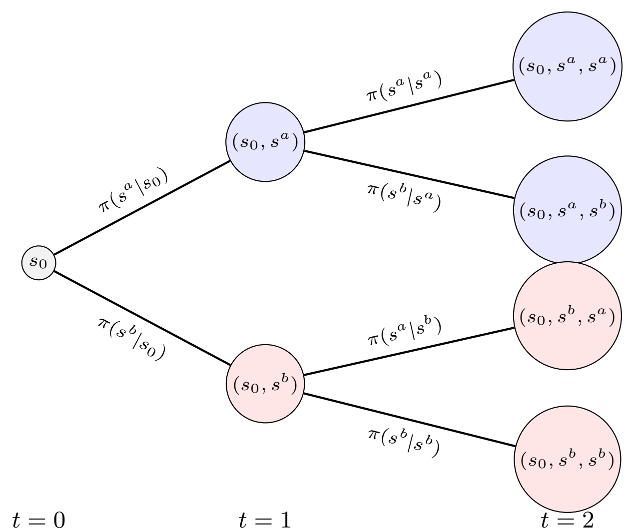

A few features deserve emphasis:

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

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

Definition 1.1: Feasible Allocation

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

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

Definition 1.2: Pareto Optimal

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

1.1 Find Pareto Optimal Allocations: Social Planner

The social planner's problem is given by

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

The Lagrange is given by

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

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

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

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

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

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

This is the condition for Pareto optimality.

Specifically,

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

And this gives

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

Feasibility constraint requires that

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

Hence we have

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

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

Remark.

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

1.2 Arrow-Debreu Trading (Time 0 Trading)

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

Definition 1.3: Competitive Equilibrium (Time-0 Trading)

A competitive equilibrium consists of

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

such that

- **Agent's Utility Max**:

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

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

- **Market Clearing**:

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

Remark.

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

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

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

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

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

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

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

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

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

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

Specifically,

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

This gives

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

From the market clearing condition, we have

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

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

Claim

A competitive equilibrium allocation is Pareto optimal when

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

Proof for Claim.

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

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

And this gives

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

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

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

And this gives

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

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

Theorem 1.4: First Welfare Theorem

A competitive equilibrium allocation is Pareto optimal.

Theorem 1.5: Second Welfare Theorem

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

Algorithm (Solve for Competitive Equilibrium)

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

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

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

3. Using budget constraint:

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

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

   Let

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

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

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

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

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

Example.

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

FOC gives

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

By the market clearing condition, we have

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

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

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

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

Plugging it back into the budget constraint, we have

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

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

Remark.

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

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

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

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

1.3 Sequential Trading

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

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

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

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

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

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

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

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

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

Definition 1.6: Natural Debt Limit

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

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

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

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

Remark.

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

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

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

Definition 1.7: Competitive Equilibrium (Sequential)

A competitive equilibrium consists of

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

such that

- **Agent's Utility Max**:

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

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

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

The Lagrangian for the agent's problem is given by

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

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

The FOC's are

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

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

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

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

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

Having this result, we can rewrite the FOC as

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

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

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

Consequently,

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

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

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

And this gives

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

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

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

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

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

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

Remark (Asset Pricing).

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

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

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

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

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

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

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

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

Define the *gross return of the asset* as

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

Define the *stochastic discount factor* as[a]

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

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

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

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

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

1.4 Recursive Competitive Equilibrium

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

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

Assumption 1.8: Markovian and Time-Invariant Fundamentals

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

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

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

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

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

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

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

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

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

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

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

Definition 1.9: Recursive Competitive Equilibrium

A recursive competitive equilibrium consists of

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

such that

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

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

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

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

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

The FOC for the Bellman equation is given by

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

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

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

The envelope condition for the Bellman equation is given by

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

Combining the conditions above, we have

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

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

1.5 Recursive Social Planner's Problem

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

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

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

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

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

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

Assume:

- Two agents with constant aggregate endowment:

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

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

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

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

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

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

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

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

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

$$\text{s.t.} \quad \sum_s \pi(s)\left(u(c_s) + \beta w_s\right) \geq v.$$

The Lagrange for the problem is given by

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

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

The FOC's are given by

- for c_s:
$$-u'(1-c_s) + \theta u'(c_s) = 0 \quad \Longrightarrow \quad \theta = \frac{u'(1-c_s)}{u'(c_s)}.$$
- for w_s:
$$\beta P'(w_s) + \theta\beta = 0 \quad \Longrightarrow \quad P'(w_s) + \theta = 0.$$

And the envelope condition is given by

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

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

From those conditions, we have

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

And we also have

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

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

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

Remark (Chapter Summary).

- **Three equivalent representations of the same equilibrium.** The complete-markets allocation can be characterized as (i) the social planner's optimum with Pareto weights λ^i, (ii) a date-0 Arrow–Debreu trading equilibrium with state-contingent prices $q(s^t)$, and (iii) a sequential-trading equilibrium with one-period-ahead Arrow securities priced by the kernel $Q(s'|s)$. The First and Second Welfare Theorems supply the bridges between them.
- **The natural debt limit $A^i(s^t)$.** Borrowing up to the present value of future endowments is the loosest possible constraint compatible with feasibility. Under Inada

conditions it never binds in equilibrium, so it can be written down without affecting the analysis.

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

Part II

Growth, Business Cycles, and Quantitative Macroeconomics

Lectures by Kai-Jie Wu

Part III

Problem Sets and Solutions

Subject Index

Author Index