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 3

Endogenously Incomplete Markets

Remark (Notation in This Chapter).

Symbol	Meaning
V^{aut}	Autarky continuation value (consume own endowment forever)
v, w_s	Promised lifetime utility today; continuation promise after state s
$P(v)$	Money lender's profit frontier in the one-sided LoC model
Δ, $\bar{\Delta}$	Surplus utility above autarky; upper bound of Δ
$Q(\Delta, s)$	Pareto frontier in the two-sided LoC model
$\bar{c}, \underline{c}$	Upper / lower bound of the consumption interval at a given state
$\lambda_s, \theta_{s'}$	Lagrange multipliers on the participation constraints
μ	Multiplier on the promise-keeping constraint
$W(s^t)$	Agent's wealth $\equiv$ PV of future endowments (Bulow–Rogoff)
$D(s^t)$	Agent's debt $\equiv$ PV of future payments (Bulow–Rogoff)
k	Debt-to-wealth limit ratio (Bulow–Rogoff)
$a(s^t), g(s^{t+1})$	Cash-in-advance saving and state-contingent payoff (Bulow–Rogoff)

3.1 One-Sided Lack of Commitment

Assume:

- Money lender (ML):
  - risk neutral.
  - can borrow and lend money at the interest rate R such that $\beta R = 1$.
  - contracts are state-dependent.
  - all borrowings and savings can only go through the money lender.
- Consumer:

– preferences:

$$\mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t u(c_t)\right],$$

where u is assumed to be strictly increasing, strictly concave, and twice continuously differentiable.

– stochastic endowment: $y(s^t)$ that is i.i.d. with each state s^t occurring with probability $\pi(s^t)$. And the space of $y(s^t)$ is finite such that

$$y_1 < y_2 < \cdots < y_N.$$

– Consumers cannot commit to a contract: if the consumer walks out of the contract at some point, then the consumer will not be able to borrow or lend from the money lender in the future. That is, from the period onwards, the consumer can only consume their endowment $y(s^t)$ in each period ("*Autarky*").

Before solving the complex model with lack of commitment, it is standard practice to establish the frictionless "First-Best" (efficient) benchmark.

Claim: First-Best Benchmark

The efficient allocation is such that the consumer consumes the average endowment at each period and each state, i.e.,

$$\sum_t \beta^t u\left(\sum_s \pi_s y_s\right) = \frac{1}{1-\beta} u\left(\sum_s \pi_s y_s\right).$$

Thus, the profit of the money lender is given by

$$\mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t \left(y(s^t) - c_t\right)\right] = \mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t \left(y(s^t) - \sum_s \pi_s y_s\right)\right].$$

Why is it efficient? The economy consists of a risk-neutral Money Lender (ML) and a strictly risk-averse consumer. A fundamental result in microeconomics is that Pareto efficiency requires the risk-neutral party to fully absorb all risk. Therefore, the ML should provide *perfect insurance*, keeping the consumer's consumption perfectly smooth across all states and time. To ensure the contract is feasible (and the ML breaks even in expectation), this constant consumption level must exactly equal the expected (average) endowment: $c_t = \sum_s \pi_s y_s := \bar{y}$.

But later we will see that this efficient allocation is not sustainable under the lack of commitment friction. Intuitively, this perfect insurance contract requires the consumer to make net payments to the ML in good states ($y_s > \bar{y}$) and receive transfers in bad states ($y_s < \bar{y}$). However, under *lack of commitment*, when a good state realizes (e.g., the best state $y_N > \bar{y}$), the consumer's autarky payoff $u(y_N)$ at the current period strictly dominates the contract's smoothed payoff $u(\bar{y})$, which makes the consumer want to walk out of the contract and consume the endowment y_N instead.

3.1.1 Self-Enforcing Contract

Due to the lack of commitment, the ML will find a contract that is *self-enforcing* (also called *sustainable*), which means that the consumer will not want to walk out of the contract at any point in time.

Assume the ML is offering a stream of consumption $\{c(s^t)\}_{t=0}^{\infty}$ to the consumer.

Under the contract, the consumer's utility is given by

$$v = \sum_t \sum_{s^t} \beta^t \pi(s^t) u(c(s^t)).$$

For the money lender, the profit is given by

$$P = \sum_t \sum_{s^t} \beta^t \pi(s^t) \left(y(s^t) - c(s^t) \right).$$

If at some point the consumer decides to walk out of the contract, then the consumer's utility will be given by

$$\begin{aligned} & u(y(s)) + \beta \sum_{s'} \pi(s'|s) u(y(s')) + \beta^2 \sum_{s''} \pi(s''|s') u(y(s'')) + \cdots \\ =& u(y(s)) + \frac{\beta}{1-\beta} \sum_{s'} \pi(s'|s) u(y(s')) \\ :=& u(y(s)) + \beta V^{\text{aut}}. \end{aligned}$$

For a contract to be self-enforcing, the ML must offer at least the utility that the consumer can get by walking out of the contract, for all t and s^t.

The natural question is then, what is the optimal contract?

The ML wants to maximize the profit P by choosing the contract $\{c(s^t)\}_{t=0}^{\infty}$ subject to the self-enforcing constraint.

To analyze this question, we make an important assumption about the timing. Specifically, we assume *ex-ante* timing, which means that at each state, the consumer wakes up and makes decisions based on the expected realizations, and then the state (and thus the endowment realization) is revealed.

The ML's problem is given by

$$\begin{aligned} P(v) = \max_{c_s, w_s} & \sum_s \pi_s \left[y(s) - c(s) + \beta P(w) \right] \\ \text{s.t.} \quad & \sum_s \pi_s \left[u(c_s) + \beta w_s \right] \geq v, \\ & u(c_s) + \beta w_s \geq u(y_s) + \beta V^{\text{aut}}, \quad \forall s, \end{aligned}$$

where the first constraint is the *promise-keeping constraint*, and the second constraint is the *participation constraint* (or, the self-enforcing constraint from the perspective of the ML).

Here we in addition assume that

$$\begin{aligned} w_s &\in \left[V^{\text{aut}}, V^{\max} \right], \quad \forall s, \\ c_s &\in \left[c_{\min}, c_{\max} \right], \quad \forall s. \end{aligned}$$

And we give the following claim without proof:

Claim

- $P(v)$ is strictly decreasing in v.
- $P(v)$ is strictly concave in v.
- $P(v)$ is continuously differentiable in v.

The Lagrange is given by

$$\begin{aligned}\mathcal{L} = &\sum_s \pi_s \left[y_s - c_s + \beta P(w_s)\right] \\ &+ \mu \left[\sum_s \pi_s \left[u(c_s) + \beta w_s\right] - v\right] \\ &+ \sum_s \lambda_s \left[u(c_s) + \beta w_s - u(y_s) - \beta V^{\text{aut}}\right].\end{aligned}$$

The FOC is given by

- w.r.t. c_s:

$$-\pi_s + \mu \pi_s u'(c_s) + \lambda_s u'(c_s) = 0.$$

- w.r.t. w_s:

$$\beta \pi_s P'(w_s) + \mu \beta \pi_s + \lambda_s \beta = 0.$$

The envelope condition is given by

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

From the three conditions above, we can obtain the following results:

- $u'(c_s) = -1/P'(w_s)$.
- $P'(v) = \lambda_s/\pi_s + P'(w_s)$.
- $1/u'(c_s) = \lambda_s/\pi_s + 1/u'(c_{s-1})$.

We analyze the problem by considering two cases:

- Case 1: Participation constraint is not binding.

  Then $\lambda_s = 0$. This immediately implies

$$P'(v) = P'(w_s), \quad 1/u'(c_s) = 1/u'(c_{s-1}).$$

  Since $P'(\cdot)$ and $u'(\cdot)$ are strictly increasing, we have $v = w_s$ and $c_s = c_{s-1}$. So when the participation constraint does not bind, the consumption and promised value are just constants and do not depend on the state s.

- Case 2: Participation constraint is binding for some s.

Then $\lambda_s > 0$. This implies

$$\begin{cases} P'(v) > P'(w_s) \implies v < w_s, \\ 1/u'(c_s) > 1/u'(c_{s-1}) \implies c_{s-1} < c_s. \end{cases}$$

Intuition: The consumer has incentives to walk out of the contract when they are in a good state (i.e., when $y(s)$ is high). To prevent the consumer from walking out of the contract, the ML needs to offer a higher promised value w_s and a higher consumption c_s in the good state than in the bad state.

In conclusion, the consumption is always non-decreasing: strictly increases when the participation constraint binds, and remains constant when the participation constraint does not bind. When the participation constraint binds (i.e., the consumer wakes up in a good state), the ML incentivizes the consumer to stay in the contract by offering a higher promised value and a higher consumption from then on (although the consumer will then need to forgo some of their current endowment since $c_s < y_s$[1]).

Motivated by the analysis above, for a given promised value v, there exists $\bar{y}(v)$ such that the participation constraint binds when $y(s) > \bar{y}(v)$ and does not bind when $y(s) < \bar{y}(v)$. In order to find such $\bar{y}(v)$, we can solve the following equations:

$$\begin{cases} w_s = v \\ c_s = c_{s-1} \\ u(c_s) + \beta w_s = u(y(s)) + \beta V^{\text{aut}}. \end{cases}$$

Below we give graphical illustrations of the above analysis. When $\lambda_s = 0$, $c_s = c_{s-1}$ and $w_s = v$. When $\lambda_s > 0$, $c_s := \tilde{c}(s)$ such that $\tilde{c}(s) > c_{s-1}$ and $w_s := \tilde{w}(s)$ such that $\tilde{w}(s) > v$.

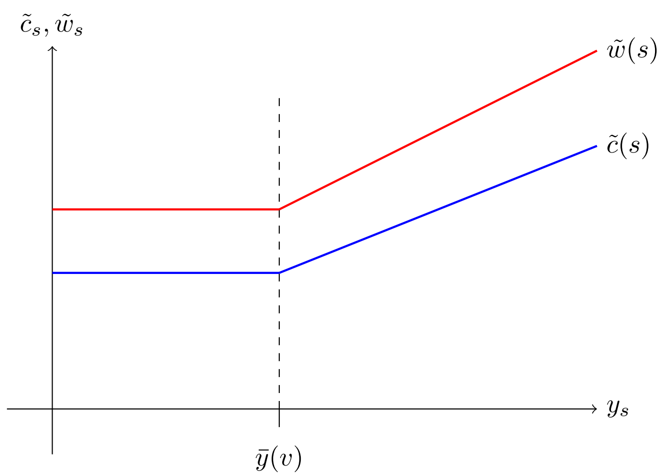

Consider the following consumption dynamics. We assume that the participation constraint does not bind at y_1, but binds at $y_2, \cdots, y_N$.

[1]Here we show that $c_s < y_s$ for all state s. In order to make the consumer participate in the contract, the ML must make the promised value $v \geq V^{\text{aut}}$. When the participation constraint binds, we have $u(c_s) + \beta w_s = u(y(s)) + \beta V^{\text{aut}}$. Since $w_s > v \geq V^{\text{aut}}$, we have $u(c_s) < u(y(s))$, which implies $c_s < y_s$. This happens when the participation constraint binds, so when the participation constraint does not bind, we still have $c_s < y_s$. Intuitively, this is true because the ML needs to make a positive profit from the contract, so the consumption stream offered to the consumer must be less than their endowment in each state.

t	1	2	3	4	5	6	7
y_t	y_1	y_1	y_2	y_1	y_N	y_2	y_1

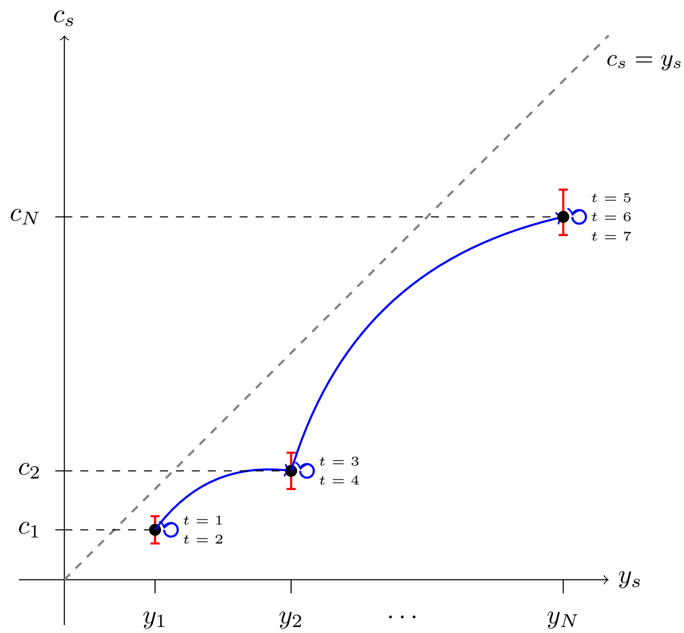

Initially, the consumer wakes up in a bad state with endowment y_1 and the participation constraint does not bind, so the consumer will consume c_1. At $t = 2$, the consumer wakes up in the same state with endowment y_1, so the consumption remains unchanged. At $t = 3$, the consumer wakes up in a good state with endowment y_2, and the participation constraint binds. In order to incentivize the consumer to stay in the contract, the ML needs to offer a higher promised value and a higher consumption c_2 such that $c_2 > c_1$. At $t = 4$, the consumer wakes up in the same state with endowment $y_1 < y_2$, so the participation constraint does not bind, and then the consumption remains unchanged at c_2. At $t = 5$, the consumer wakes up in a better state with endowment y_N, and the participation constraint binds again. The ML then offers a higher promised value and a higher consumption c_N such that $c_N > c_2$ to keep the consumer in the contract. At $t = 6$ and 7, although the consumer wakes up in relatively bad states, the consumption remains unchanged (in particular, does not decrease) at c_N since the participation constraint does not bind.

3.1.2 Profit of the ML

Ex-ante, the ML would offer

$$v = V^{\text{aut}}.$$

If y_s happens to be y_1, and $v = V^{\text{aut}}$, we consider the case when the participation constraint binds:

$$u(c_1) + \beta v = u(y_1) + \beta V^{\text{aut}}.$$

This implies the ML must offer the consumer $c_1 = y_1$ to make them stay in the contract. Then the ML will make zero profit from the consumer when $y_s = y_1$. (From the reasoning we also know that $\bar{y}(V^{\text{aut}}) = y_1$.)

When y_s happens to be y_N at some point, the consumption will be c_N from then on. In other words, the consumption will converge to $c_N := \bar{c}$ once the best state y_N is realized.

Again, the ML will offer c_N to make the consumer just stay in the contract when $y_s = y_N$:

$$u(\bar{c}) + \beta\frac{u(\bar{c})}{1-\beta} = u(y_N) + \beta V^{\text{aut}} = u(y_N) + \frac{\beta}{1-\beta}\sum_s \pi_s u(y_s).$$

Previously we have shown that $c_s < y_s$ for all state s, so we have $\bar{c} < y_N$. This implies that

$$u(\bar{c}) > \sum_s \pi_s u(y_s).$$

Namely, the value of the contract is higher than the (expected) value of autarky.

Once the best state y_N is reached, the profit of the ML is given by

$$\begin{aligned} P &= y_N - \bar{c} + \frac{\beta}{1-\beta}\sum_s \pi_s(y_s - \bar{c}) \\ &= y_N + \frac{\beta}{1-\beta}\sum_s \pi_s y_s - \frac{\bar{c}}{1-\beta}. \end{aligned}$$

Claim

The profit of the ML is positive after the best state y_N is reached.

Proof for Claim.

The profit of the ML is positive if and only if

$$y_N + \frac{\beta}{1-\beta}\sum_s \pi_s y_s > \frac{\bar{c}}{1-\beta}.$$

We prove this by contradiction. Suppose not, then we have

$$\begin{aligned} & y_N + \frac{\beta}{1-\beta}\sum_s \pi_s y_s \le \frac{\bar{c}}{1-\beta} \\ \iff & (1-\beta)y_N + \beta\sum_s \pi_s y_s \le \bar{c} \\ \iff & u\left((1-\beta)y_N + \beta\sum_s \pi_s y_s\right) \le u(\bar{c}) = (1-\beta)u(y_N) + \beta\sum_s \pi_s u(y_s), \\ \iff & u\left((1-\beta)y_N + \beta\sum_s \pi_s y_s\right) \le (1-\beta)u(y_N) + \beta\sum_s \pi_s u(y_s). \end{aligned}$$

where the equality in the third line holds from the participation constraint:

$$u(\bar{c}) + \beta\frac{u(\bar{c})}{1-\beta} = u(y_N) + \frac{\beta}{1-\beta}\sum_s \pi_s u(y_s) \implies u(\bar{c}) = (1-\beta)u(y_N) + \beta\sum_s \pi_s u(y_s).$$

But the last inequality holds if and only if u is convex, by Jensen's inequality. This contradicts the assumption that u is strictly concave. Therefore, we must have $P > 0$

after the best state y_N is reached.

Remark (Intuition).

- Intuitively, the ML is making positive profit because of risk sharing. The consumer is risk averse, while the ML is risk neutral. After hitting the best state, the consumer is insured against the idiosyncratic risk through the contract, so the consumer is willing to pay a positive amount (by sacrificing the current consumption of $y_N - c_N > 0$) to the ML to get insurance. The ML then makes positive profit from the contract (*risk premium*).

- The agent owes money to the ML in y_N state, but receives net transfers (insurance) from the ML in future bad states. The ML makes positive profit from the contract because the consumer is risk averse and thus values the insurance provided by the contract.

Remark (Why focus only on y_1 and y_N? What about the intermediate states?).

In the analysis of the one-sided lack of commitment model, the notes explicitly target y_1 and y_N because they serve as the two critical **boundary conditions** of the dynamic contract. The intermediate states are omitted not because they do not exist, but because characterizing the boundaries is sufficient to demonstrate the economic mechanism of the contract.

In the intermediate states, the single-period profit $y_s - c_s$ is transitional and highly history-dependent. The ML might suffer temporary losses (paying out insurance when the agent draws bad shocks) or earn small profits. To establish that the contract is self-enforcing and profitable *ex-ante*, however, it suffices to show that the ML breaks even at the bottom of the state space (y_1) and strictly collects a risk premium once the top (y_N) is reached. The intermediate states then take care of themselves.

3.2 Bulow-Rogoff Contracts

This model deviates from the standard one-sided lack of commitment model. In standard models (e.g., Eaton-Gersovitz 1981), the punishment for default is a permanent exclusion from future borrowing, which acts as a sufficient threat to sustain debt. This model introduces a critical loophole: it allows the agent to ***save*** *in complete financial markets after defaulting. We will prove that if an agent has the ability to save after default, borrowing through the Money Lender (ML) becomes completely infeasible (i.e., the credit limit inevitably collapses to zero).*

We assume

- **Agent:** Risk-averse with strictly increasing utility $u(\cdot)$. Receives a stochastic endowment stream $y(s)$.

- **Money Lender (ML):** Risk-neutral and competitive. Has access to complete financial markets and discounts the future at a risk-free gross interest rate $R > 1$.

Definition 3.1: Wealth

The agent's wealth at node s^t, denoted by $W(s^t)$, is defined as the expected present discounted value of all future endowments:

$$W(s^t) = \sum_{\tau \geq t} \sum_{s^\tau | s^t} \pi(s^\tau | s^t) \frac{y(s^\tau)}{R^{\tau - t}}$$

In addition, we assume $W(s^t)$ is finite (which requires R to be sufficiently high).

Definition 3.2: Contract

A contract is defined by a stream of payments $p(s^t)$ from the agent to the ML.

- $p(s^t) > 0$: The agent is paying back the ML.
- $p(s^t) < 0$: The agent is borrowing (receiving inflows from the ML).

Definition 3.3: Debt

The agent's debt at node s^t, denoted by $D(s^t)$, is defined as the expected present discounted value of all future payments:

$$D(s^t) = \sum_{\tau \geq t} \sum_{s^\tau | s^t} \pi(s^\tau | s^t) \frac{p(s^\tau)}{R^{\tau - t}}$$

Intuitively, the ML manages risk by enforcing a strict credit limit: the agent's debt can never exceed a specific fraction k of their total future wealth.

Assumption 3.4: Borrowing Limit

The ML will define a parameter $k \in [0, 1]$ such that:

$$D(s^t) \leq kW(s^t), \quad \forall s^t.$$

In default, the agent can no longer borrow from the ML, but can still save by purchasing cash-in-advance contracts (Arrow securities).

Definition 3.5: Cash-In-Advance

Cash-in-advance contracts allow the agent to save by purchasing state-contingent claims on future payoffs. The agent can choose a saving strategy defined by:

$$a(s^t) = \sum_{s^{t+1} | s^t} \pi(s^{t+1} | s^t) \frac{g(s^{t+1})}{R}$$

where $g(s^{t+1}) \geq 0$ is the state-contingent return on the asset tomorrow.

The intuition if the cash-in-advance contracts is that the agent acts as their own actuary. $a(s^t)$ is the total premium (cash) the agent pays today. In exchange, the market guarantees a non-negative payout $g(s^{t+1})$ tomorrow, exactly tailored to the agent's desired self-insurance needs.

Theorem 3.6: Bulow-Rogoff

In any equilibrium, the debt must be non-positive ($D(s^t) \leq 0$) for all s^t, the borrowing limit $k = 0$, and the agent cannot borrow at all.

Here, we motivate the proof by intuition before going into the formal proof.

Suppose there exists a node s^t where the debt satisfies two conditions:

1. $D(s^t) \leq kW(s^t)$ (The global borrowing limit is respected)
2. $D(s^t) > k\left(W(s^t) - y(s^t)\right)$ (The debt exceeds future collateral)

The intuition for the second condition is that, the current debt is so heavy that it exceeds the allowable borrowing limit on *strictly future* wealth $(W - y)$. Consequently, the agent is forced to make a painful net transfer out of today's pocket $(y(s^t))$ to service the debt. This is the optimal moment to default.

We consider the case when the agent defaults at this specific node s^t and starts the following savings sequence for all $\tau \geq t$:

$$a(s^\tau) = p(s^\tau) + k\left(W(s^\tau) - y(s^\tau)\right) - D(s^\tau),$$

where k is assumed to be $k > 0$, since we try to prove by contradiction that k must be zero in equilibrium.

Here is the intuition for this saving sequence: the agent redirects the payment $p(s^\tau)$ they *would* have made to the ML into a private savings account. They adjust this amount by taking the maximum theoretical future borrowing limit $(k(W - y))$ and subtracting the actual future debt burden (D). In effect, they use the financial market to mimic the ML's balance sheet.

But we are still left with two critical questions:

- **Feasibility:** Is this asset (defined by the savings sequence) legally available in the market? In other words, can the agent actually purchase this asset to implement the savings strategy?

- **Optimality:** Does this asset make the agent strictly better off than repaying the debt? In other words, is defaulting and following this savings strategy strictly more profitable than repaying the debt?

To prove this asset is legally available in the market, we must show that its future payoff $g(s^{\tau+1})$ is non-negative.

First, we write down the recursive definitions of W and D:

$$W(s^\tau) = y(s^\tau) + \sum_{s^{\tau+1}|s^\tau} \pi(s^{\tau+1}|s^\tau)\frac{W(s^{\tau+1})}{R} \implies W(s^\tau) - y(s^\tau) = \sum_{s^{\tau+1}|s^\tau} \pi(s^{\tau+1}|s^\tau)\frac{W(s^{\tau+1})}{R}$$

$$D(s^\tau) = p(s^\tau) + \sum_{s^{\tau+1}|s^\tau} \pi(s^{\tau+1}|s^\tau)\frac{D(s^{\tau+1})}{R} \implies D(s^\tau) - p(s^\tau) = \sum_{s^{\tau+1}|s^\tau} \pi(s^{\tau+1}|s^\tau)\frac{D(s^{\tau+1})}{R}$$

Substitute these recursive forms back into our savings sequence $a(s^\tau)$:

$$\begin{aligned}
a(s^\tau) &= k\left(W(s^\tau) - y(s^\tau)\right) - \left(D(s^\tau) - p(s^\tau)\right) \\
&= k \sum_{s^{\tau+1}|s^\tau} \pi(s^{\tau+1}|s^\tau) \frac{W(s^{\tau+1})}{R} - \sum_{s^{\tau+1}|s^\tau} \pi(s^{\tau+1}|s^\tau) \frac{D(s^{\tau+1})}{R} \\
&= \sum_{s^{\tau+1}|s^\tau} \pi(s^{\tau+1}|s^\tau) \frac{kW(s^{\tau+1}) - D(s^{\tau+1})}{R}
\end{aligned}$$

By matching this with the standard cash-in-advance pricing formula ($a(s^\tau) = \sum \pi \frac{g}{R}$), we identify the state-contingent payoff tomorrow as:

$$g(s^{\tau+1}) = kW(s^{\tau+1}) - D(s^{\tau+1}).$$

Since the ML's contract globally enforces $D(s^{\tau+1}) \leq kW(s^{\tau+1})$, it strictly follows that $g(s^{\tau+1}) \geq 0$ for all possible future states. Therefore, the proposed asset is a feasible cash-in-advance contract.

Then, we compare the agent's consumption payoffs under repayment vs. default.

- At the moment of default (Today, s^t):
  - Payoff in repayment: $y(s^t) - p(s^t)$
  - Payoff in default:

$$\begin{aligned}
y(s^t) - a(s^t) &= y(s^t) - \left[p(s^t) + k\left(W(s^t) - y(s^t)\right) - D(s^t)\right] \\
&= y(s^t) - p(s^t) + \underbrace{\left[D(s^t) - k\left(W(s^t) - y(s^t)\right)\right]}_{>0 \text{ (due to our choice of the tipping point)}}
\end{aligned}$$

  As a result, the agent consumes strictly more today by defaulting.

- At any future node (Tomorrow and beyond, s^τ for $\tau > t$):
  - Payoff in repayment: $y(s^\tau) - p(s^\tau)$
  - Payoff in default (= endowment - cost of new asset + payoff from yesterday's asset):

$$\begin{aligned}
&y(s^\tau) - a(s^\tau) + g(s^\tau) \\
&= y(s^\tau) - [p(s^\tau) + k(W(s^\tau) - y(s^\tau)) - D(s^\tau)] + [kW(s^\tau) - D(s^\tau)] \\
&= y(s^\tau) - p(s^\tau) - kW(s^\tau) + ky(s^\tau) + D(s^\tau) + kW(s^\tau) - D(s^\tau) \\
&= (1+k)y(s^\tau) - p(s^\tau)
\end{aligned}$$

  As a result, since the endowment $y > 0$, if $k > 0$, then $(1+k)y - p > y - p$. The agent consumes strictly more (ky extra) in every single future period.

In conclusion, if $k > 0$, the agent will inevitably reach node s^t, choose to default, and be strictly better off in every period thereafter. Anticipating this, the ML will never lend. Consequently, the only equilibrium is $k \leq 0 \implies k = 0$, meaning debt $D(s^t) \leq 0$, and the agent cannot borrow.

Remark (Existence of Tipping Point and Discussion of k).

The proof above relies on the existence of a **tipping point** s^* that satisfies the two conditions. Existence is guaranteed by the following lemma.

Lemma 3.7: Existence of the Tipping Point

Suppose the Money Lender (ML) offers a contract where the supremum of the debt-to-wealth ratio is strictly positive. Let $k = \sup_{s^t} \frac{D(s^t)}{W(s^t)} > 0$. Then, there exists at least one node s^* in the contract history that satisfies:

1. $D(s^*) \leq kW(s^*)$
2. $D(s^*) > k\left(W(s^*) - y(s^*)\right)$

Proof for Lemma

By the definition of the supremum k, the first condition $D(s^t) \leq kW(s^t)$ holds globally for all s^t.

To prove the second condition, we first establish a strictly positive lower bound for the ratio of the flow endowment to the total stock wealth, $\frac{y(s^t)}{W(s^t)}$. Assuming a finite Markov state space where endowments are strictly positive, there exists a minimum possible endowment $y_{\min} > 0$ and a maximum possible total wealth $W_{\max} < \infty$ (since the discount rate $R > 1$). Therefore, for any node s^t:

$$\frac{y(s^t)}{W(s^t)} \geq \frac{y_{\min}}{W_{\max}} \equiv \delta > 0$$

Now, choose an arbitrarily small positive number ϵ such that $\epsilon = \frac{1}{2}k\delta > 0$. By the definition of the supremum k, there must exist at least one node s^* such that the actual debt-to-wealth ratio is strictly within ϵ of the supremum:

$$\frac{D(s^*)}{W(s^*)} > k - \epsilon \implies D(s^*) > kW(s^*) - \epsilon W(s^*)$$

Substitute our specific choice of ϵ into this inequality:

$$\begin{aligned} D(s^*) &> kW(s^*) - \left(\frac{1}{2}k\delta\right) W(s^*) \\ &> kW(s^*) - k\left(\frac{y(s^*)}{W(s^*)}\right) W(s^*) \\ &= k\left(W(s^*) - y(s^*)\right). \end{aligned}$$

Thus, the node s^* satisfying both conditions is mathematically guaranteed to exist. ■

For narrative simplicity, we often state that the Money Lender (ML) specifies a global borrowing limit k such that $D(s^t) \leq kW(s^t)$. However, mathematically, it is crucial that the k used in this proof is defined as the *supremum* of the actual debt-to-wealth ratio, rather than an arbitrary constant.

If we were to use a loosely defined nominal rule k_{rule}, and the actual borrowing behavior is significantly more conservative (i.e., $k_{\text{rule}} \gg \sup \frac{D}{W}$), the strict inequality

$D(s^t) > k_{\text{rule}} \left(W(s^t) - y(s^t)\right)$ might never be satisfied. In such a case, the debt would never mathematically "squeeze" the current endowment enough to trigger the tipping point. Therefore, constructing the shadow asset around the precise, tight supremum k is the exact mechanism that guarantees the existence of the default node.

3.3 Two-Sided Lack of Commitment

Assume:

- Two agents:
  - identical preferences with risk aversion;
  - lack of commitment.
- Endowments:

  $$y_A = y_s, \quad y_B = 1 - y_s, \quad y_s \in [0, 1],$$

  where y_s is i.i.d. with each state s occurring with probability $\pi(s)$.
  - The feasibility constraint is then given by

  $$c_A(s^t) + c_B(s^t) = 1, \quad \forall t, s^t.$$

- Lending and borrowing can only be done through the reallocation of consumption between agents. If any agent defaults on the contract, then both agents will be in autarky and consume their own endowment forever.

$$V_A^{\text{aut}} = \sum_t \beta^t \sum_s \pi(s) u(y_s) = \frac{\sum_s \pi_s u(y_s)}{1 - \beta},$$
$$V_B^{\text{aut}} = \sum_t \beta^t \sum_s \pi(s) u(1 - y_s) = \frac{\sum_s \pi_s u(1 - y_s)}{1 - \beta}.$$

An allocation $\{c_A(s^t), c_B(s^t)\}_{t=0}^{\infty}$ is sustainable if neither agent has incentives to default on the contract at any point in time. That is, for all t and s^t, we have

$$u(c^A(s^t)) + \mathbb{E}_t \left[\sum_j \beta^j u(c^A(s^{t+j})) \right] \geq u(y(s^t)) + \beta V_A^{\text{aut}},$$
$$u(c^B(s^t)) + \mathbb{E}_t \left[\sum_j \beta^j u(c^B(s^{t+j})) \right] \geq u(1 - y(s^t)) + \beta V_B^{\text{aut}}.$$

We characterize the Pareto frontier of feasible and sustainable allocation by the function $Q(v)$, which is defined as the maximum utility that agent B can get given that agent A gets at least v more utility than their autarky utility (we call it *difference in utilities* later):

$$
\begin{aligned}
Q(v) = \max_{c(s^t)} & \sum_t \sum_{s^t} \beta \left[u(1-c(s^t)) - u(1-y(s^t))\right] \\
\text{s.t.} \quad & \sum_t \sum_{s^t} \beta \left[u(c(s^t)) - u(y(s^t))\right] \geq v, \\
& c(s^t) \in \Gamma.
\end{aligned}
$$

where Γ is the set of feasible and sustainable allocations.

Obviously, it has to be true that $v \geq 0$ and $Q(v) \geq 0$ for all $v \geq 0$.

If $\underline{v} = 0$, then the present value of A's consumption utility is promised to be just at least as high as A's autarky utility, so $v_{\min} = \underline{v} = 0$. If $Q(\bar{v}) = 0$, then the present value of B's consumption utility is promised to be just at least as high as B's autarky utility, so $v_{\max} = \bar{v}$.

The recursive problem can be written as

$$
\begin{aligned}
Q(\Delta, s) = \max_{c, \Delta(s')} \ & u(1-c) - u(1-y_s) + \beta \sum_{s'} \pi(s') Q(\Delta(s'), s') \\
\text{s.t.} \ & u(c) - u(y_s) + \beta \sum_{s'} \pi(s') \Delta(s') \geq \Delta, \\
& \Delta(s') \geq 0, \quad \forall s', \\
& Q(\Delta(s'), s') \geq 0, \quad \forall s', \\
& c \in [0, 1].
\end{aligned}
$$

Here Δ is the promise in terms of the difference in utilities for agent A (i.e., the present value of A's consumption utility minus A's autarky utility). Note that the first constraint is the promise-keeping constraint, and the second and third constraints are the participation constraints for agent A and agent B, respectively.

Remark.

- We can show that $\Delta(s) \in \left[0, \bar{\Delta}(s)\right]$ for all s, where $\bar{\Delta}(s)$ is such that $Q(\bar{\Delta}(s), s) = 0$.
- Q inherits the properties of u: Q is decreasing in Δ, strictly concave in Δ, and continuously differentiable in Δ.

Let μ be the Lagrange multiplier for the promise-keeping constraint, $\lambda_{s'} \beta \pi_{s'}$ be the Lagrange multiplier for the participation constraint of agent A, and $\theta_{s'} \beta \pi_{s'}$ be the Lagrange multiplier for the participation constraint of agent B. The FOC is given by

- w.r.t. c:

$$-u'(1-c) + \mu u'(c) = 0.$$

- w.r.t. $\Delta(s')$:

$$\beta \pi_{s'} Q'(\Delta(s'), s') + \mu \beta \pi_{s'} + \lambda_{s'} \beta \pi_{s'} + \theta_{s'} \beta \pi_{s'} Q'(\Delta(s'), s') = 0.$$

The envelope condition is given by

$$Q'(\Delta, s) = -\mu.$$

From the FOC and the envelope condition, we can obtain the following results:

$$Q'(\Delta, s) = -\frac{u'(1-c)}{u'(c)},$$
$$Q'(\Delta, s) = (1+\theta_{s'})Q'(\Delta(s'), s') + \lambda_{s'}.$$

This implies that the consumption c is increasing in Δ. We have already argued that $\Delta(s) \in \left[0, \bar{\Delta}(s)\right]$ for all s. Therefore, the consumption c will fall in the range $[\underline{c}_s, \bar{c}_s]$ for all s, where

$$Q'(0, s) = -\frac{u'(1-\underline{c}_s)}{u'(\bar{c}_s)}, \quad (c = \underline{c}_s, \Delta = 0),$$
$$Q'(\bar{\Delta}(s), s) = -\frac{u'(1-\bar{c}_s)}{u'(\underline{c}_s)}, \quad (c = \bar{c}_s, \Delta = \bar{\Delta}(s)).$$

Furthermore, we can rewrite the FOC w.r.t. c as

$$c = g(Q'(\Delta, s)) = g\left(-\frac{u'(1-c)}{u'(c)}\right),$$

where g^{-1} is decreasing and negative.

With this, the FOC w.r.t. $\Delta(s')$ can be rewritten as

$$g^{-1}(c_s) = (1+\theta_{s'})g^{-1}(c_{s'}) + \lambda_{s'}.$$

Here we analyze the problem by considering four cases:

- No participation constraint binds. Then $\lambda_{s'} = 0$ and $\theta_{s'} = 0$ for all s'. This implies that $Q'(\Delta, s) = Q'(\Delta(s'), s')$ and $g^{-1}(c_s) = g^{-1}(c_{s'})$ for all s and s', so the consumption is constant across states:
$$c_s = c_{s'}, \quad \Delta = \Delta(s').$$

- Only participation constraint of agent A binds. Then $\lambda_{s'} > 0$ and $\theta_{s'} = 0$ for some s'. Also by construction, $\Delta(s') = 0$ and $c_{s'} = \underline{c}_{s'}$ in this case. So we have
$$g^{-1}(c_s) = g^{-1}(c_{s'}) + \lambda_{s'}$$
$$\Longrightarrow g^{-1}(c_s) > g^{-1}(c_{s'})$$
$$\Longrightarrow c_s < c_{s'}.$$

- Only participation constraint of agent B binds. Then $\lambda_{s'} = 0$ and $\theta_{s'} > 0$ for some s'. Also by construction, $\Delta(s') = \bar{\Delta}(s')$, $Q(\Delta(s'), s') = 0$, and $c_{s'} = \bar{c}_{s'}$ in this case. So we have
$$g^{-1}(c_s) = (1+\theta_{s'})g^{-1}(c_{s'})$$
$$\Longrightarrow g^{-1}(c_s) < g^{-1}(c_{s'})$$
$$\Longrightarrow c_s > c_{s'}$$
$$\Longrightarrow 1 - c_s < 1 - c_{s'}.$$

- Both participation constraints bind.

  This case is not feasible as long as some risk-sharing is possible. If both participation constraints bind simultaneously in some state s', it implies that $\Delta(s') = 0$ (Agent A gets exactly their autarky value) and $Q(\Delta(s'), s') = 0$ (Agent B gets exactly their autarky value). Mathematically, this means $Q(0, s') = 0$.

  However, this contradicts the assumption that some risk-sharing is possible. If risk-sharing is sustainable, the contract generates a strictly positive total surplus relative to autarky due to the agents' risk aversion.

  Since $Q(\Delta, s')$ defines the *Pareto frontier* (the maximum possible surplus for B given A's surplus), if A is pushed down to their autarky outside option ($\Delta = 0$), Agent B must capture *all* the strictly positive risk-sharing surplus. Therefore, it must be true that $Q(0, s') > 0$.

  In geometric terms, the autarky point $(0, 0)$ lies strictly *inside* the feasible payoff set. The Pareto frontier Q strictly bounds this set away from the origin. The only scenario where $Q(0, s') = 0$ holds is when the commitment friction is so severe that the risk-sharing market completely collapses, making autarky the only sustainable allocation.

In conclusion,

$$c_{s'} = \begin{cases} c_s & \text{if } c_s \in [\underline{c}_{s'}, \bar{c}_{s'}] \quad \text{(no participation constraint binds)} \\ \underline{c}_{s'} & \text{if } c_s < \underline{c}_{s'}, \Delta(s') = 0 \quad \text{(only participation constraint of agent } A \text{ binds)} \\ \bar{c}_{s'} & \text{if } c_s > \bar{c}_{s'}, Q(\Delta(s'), s') = 0 \quad \text{(only participation constraint of agent } B \text{ binds)} \end{cases}$$

Remark.

- **Intuition**: When the participation constraint of agent A binds, the planner raises the consumption of agent A so that they are indifferent between autarky and the contract, given their high endowment today. The same logic applies to agent B when the participation constraint of agent B binds.

- Note that although agent A is getting the minimum consumption level $\underline{c}_{s'}$ when the participation constraint of agent A binds, it does not necessarily mean that agent A is worse off than in autarky. Actually agent A is made (weakly) better off than in the previous state because the planner raises their consumption level to keep them in the contract. (The "worst" in a good state is still better than some consumption level in a bad state.)

Claim

We can obtain a stationary distribution of consumption if

- consumption intervals do not contain each other. That is,

$$y_1 > y_2 \implies \bar{c}_1 > \bar{c}_2, \quad \underline{c}_1 > \underline{c}_2.$$

- endowments are contained in the consumption intervals. That is,

$$y_s \in [\underline{c}_s, \bar{c}_s], \quad \forall s.$$

- consumption intervals are non-degenerate. That is,

$$\underline{c}_s < \bar{c}_s, \quad \forall s.$$

Proof for Claim.

Claim

$$Q(\Delta + u(y_2) - u(y_1), s_1) = Q(\Delta, s_2) + u(1 - y_2) - u(1 - y_1).$$

Proof for Claim.

We prove this by definition.

$$\begin{aligned} Q(\Delta + u(y_2) - u(y_1), s_1) = \max u(1-c) - u(1-y_1) + \beta \sum_{s'} \pi(s') Q(\Delta(s'), s') \\ \text{s.t. } u(c) - u(y_1) + \beta \sum_{s'} \pi(s')\Delta(s') \geq \Delta + u(y_2) - u(y_1). \end{aligned}$$

$$\begin{aligned} Q(\Delta, s_2) + u(1-y_2) - u(1-y_1) = \max u(1-c) - u(1-y_2) + \beta \sum_{s'} \pi(s') Q(\Delta(s'), s') \\ \text{s.t. } u(c) - u(y_2) + \beta \sum_{s'} \pi(s')\Delta(s') \geq \Delta. \end{aligned}$$

Hence, the equality holds. ■

Let $y_1 > y_2$. We evaluate at $\Delta = \bar{\Delta}(s_2) := \bar{\Delta}_2$. By the just-proven equality, we have

$$Q(\bar{\Delta}_2 + u(y_2) - u(y_1), s_1) = Q(\bar{\Delta}_2, s_2) + u(1 - y_2) - u(1 - y_1).$$

Note that $Q(\bar{\Delta}_2, s_2) = 0$ by definition, and $u(1-y_2) - u(1-y_1) > 0$ since $y_1 > y_2$ and u is increasing. Therefore, we have $Q(\bar{\Delta}_2) > 0$. Also note that $0 = Q(\bar{\Delta}_1, s_1)$, so we have

$$Q(\bar{\Delta}_2 + u(y_2) - u(y_1), s_1) > Q(\bar{\Delta}_1, s_1).$$

Since Q is decreasing in Δ, this implies

$$\bar{\Delta}_2 + u(y_2) - u(y_1) < \bar{\Delta}_1.$$

Since Q is strictly concave, Q' is strictly decreasing in Δ. Therefore, we have

$$\begin{aligned} Q'(\bar{\Delta}_1, s_1) &< Q'(\bar{\Delta}_2 + u(y_2) - u(y_1), s_1) \\ &= Q'(\bar{\Delta}_2, s_2), \end{aligned}$$

where the equality holds by the just-proven equality (and taking derivatives).

Note that we previously defined $c = g(Q'(\Delta, s))$, and we have argued that g is strictly increasing. Finally, we have

$$\bar{c}_1 > \bar{c}_2.$$

Remark (Endogenously Incomplete Markets: Where We Have Arrived).

The three settings developed in this chapter—one-sided lack of commitment, Bulow–Rogoff with re-access to savings markets, and two-sided lack of commitment—share a common methodological signature: rather than *assuming* that the asset market is incomplete (as in the previous chapter), they let the asset market *become* incomplete as the equilibrium response to a participation friction. Three lessons survive across the three models.

- **The participation constraint is a state variable in disguise.** Whether expressed as a sequence of $u(c) + \beta w \geq u(y) + \beta V^{\text{aut}}$ inequalities (one-sided), as a wealth-to-debt ratio bound (Bulow–Rogoff), or as a pair of value-function constraints (two-sided), the shadow value of relaxing the participation constraint enters the dynamics directly. Promised utility v (or Δ) is the natural recursive state because it summarizes everything the planner needs to remember.

- **Risk sharing is partial and history-dependent.** Optimal contracts implement *conditional* consumption smoothing: consumption is constant within "intervals" of states (where no participation constraint binds) and adjusts only when a constraint binds. The result is a non-trivial stationary distribution of consumption—qualitatively similar to what exogenously incomplete markets produce, but micro-founded by the friction.

- **What outside option is available matters enormously.** Bulow–Rogoff is the cleanest illustration: relative to one-sided LoC, the only change is that the defaulting agent is allowed to save in complete markets afterward, and the entire borrowing market collapses. The lesson generalizes: any structural model of credit-market frictions stands or falls with its specification of what the agent can do after default.

These insights motivate the heterogeneous-agent macro models we turn to in later chapters: the cross-sectional wealth distribution is not just a residual of model-driven heterogeneity in shocks, but the equilibrium outcome of microfounded participation and information frictions.

Remark (Chapter Summary).

- **Endogenously incomplete markets.** Rather than restricting the asset menu by fiat, this chapter derives incompleteness as the equilibrium response to a participation friction. The set of sustainable contracts is shaped by what each party can do upon walking away.

- **Promised utility as the recursive state.** In the one-sided LoC model, v summarizes the lender's outstanding obligation. In the two-sided model, Δ (utility surplus over autarky) plays the same role. The Pareto frontiers $P(v)$ and $Q(\Delta, s)$ encode the constrained-efficient allocations.

- **Contracts smooth conditionally.** Consumption is constant within "intervals" of states (no participation constraint binds) and adjusts only when one binds. The result is an endogenous, history-dependent stationary consumption distribution—qualitatively similar to the exogenously incomplete case but micro-founded.

- **Bulow–Rogoff: the outside option is everything.** Allow defaulters to save in complete cash-in-advance markets, and the entire borrowing market collapses ($k = 0$ in equilibrium). The lesson: any structural model of credit-market frictions stands or falls with its specification of post-default options.

- **Two-sided LoC stationary distribution.** Existence of a non-degenerate stationary distribution of consumption requires the consumption intervals not to nest—i.e., risk sharing is partial but persistent.

- **Connection to Part II.** The wealth distributions that drive HANK and Aiyagari-style models can be interpreted as reduced forms of these participation-friction equilibria, justifying their economic content.

Part II

Growth, Business Cycles, and Quantitative Macroeconomics

Lectures by Kai-Jie Wu

Part III

Problem Sets and Solutions

Subject Index

Author Index