Preface

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

Audience and Prerequisites

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

Structure of the Book

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

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

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

Pedagogical Conventions

Several typographic conventions recur throughout the text.

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

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

Reading Paths

Readers do not have to proceed linearly.

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

Acknowledgments

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

Rui Zhou, Spring 2026

Notation

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

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

Symbols used throughout the book.

Symbol	Meaning
Preferences and discounting	
$u(\cdot)$	Period utility function; $u' > 0$, $u'' < 0$, satisfying Inada conditions where needed.
β	Time discount factor; $\beta \in (0, 1)$.
σ	Coefficient of relative risk aversion under CRRA utility; the inverse $1/\sigma$ is the intertemporal elasticity of substitution.
γ	Coefficient of *absolute* risk aversion under CARA utility (Ch. 2 only).
$\mathbb{E}_t[\cdot]$	Expectation conditional on history s^t.
Stochastic environment	
s_t, s^t	Date-t state; history through t.
$\pi(s^t)$	Unconditional probability of history s^t; $\pi(s^\tau \| s^t)$ is conditional.
ε_t	Innovation / shock realization.
ρ	Persistence parameter of an AR(1) process; $\rho = \psi$ in Ch. 2's CARA example.
Endowment and production	
$y(s^t)$, Y_t	Stochastic endowment; aggregate output.

(continued on next page)

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

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

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

Part I

Heterogeneous Agents in Complete and Incomplete Markets

Lectures by Maria-Jose Carreras-Valle

Part II

Growth, Business Cycles, and Quantitative Macroeconomics

Lectures by Kai-Jie Wu

Chapter 1

Consumption and Saving

Remark (Notation in This Chapter).

Symbol	Meaning
Y_t	Period-t income (deterministic in PIH section, stochastic in RWH section)
A_t	Period-t asset holdings
PI	Permanent income $\equiv A_0 + \sum_{t\geq 0} Y_t/(1+r)^t$
$\phi(r)$	Annuity factor, $\phi(r) = (1 - 1/(1+r))/(1 - (1/(1+r))^{T+1})$
Y_i^P, Y_i^T	Permanent and transitory components of cross-sectional income
$\hat{b}$	Cross-sectional regression coefficient (Keynes), $= \text{Var}(Y^P)/\text{Var}(Y)$
σ	Coefficient of relative risk aversion under CRRA preferences
$X_t = \beta^t u'(C_t)$	Discounted marginal utility (martingale under RWH)
ε_{t+1}	Consumption innovation $C_{t+1} - C_t$ in the Hall test
λ	Hand-to-mouth fraction in Campbell–Mankiw
PFD	Alaska Permanent Fund Dividend
$\hat{\alpha}$	Excess-sensitivity coefficient (Hsieh, Kueng)
s^t	History of income realizations through t (RWH stochastic setup)

The previous chapters spent considerable effort *computing* aggregate consumption and saving in stochastic, general-equilibrium models. We now step back and ask the underlying behavioral question: **how does a single household decide how much to consume in each period given a stream of income?** The answer that emerges—once we take optimization and forward-looking behavior seriously—is sharper and more counterintuitive than the simple rule "consume what you earn." This is the core insight of the **Permanent Income Hypothesis (PIH)** of **?** and ?, together with its stochastic refinement, **?**'s (?) **Random Walk Hypothesis (RWH)**.

We develop the theory in two stages. The first half of the chapter treats the household problem under *certainty*, deriving the PIH in its cleanest form. The second half reintroduces uncertainty and derives Hall's random-walk implication. The final section previews the modern empirical literature that tests both.

1.1 The Big Picture: Consumption & Saving, Data vs. Theory

The long arc of macro consumption theory pits two views against one another:

- **?:** consumption is mainly a function of *current* income. The marginal propensity to consume (MPC) out of an extra dollar today is positive and stable, around 0.6–0.8.
- **? PIH:** consumption is a function of *permanent income*—the annuity value of total lifetime resources. Transitory income shocks barely affect consumption; only permanent shifts matter.

These two views make sharply different predictions about how households respond to a tax cut, a year-end bonus, a winning lottery ticket, or a temporary recession. Most of the modern consumption literature can be read as adjudicating between them, modifying both, or measuring how far reality lies from each.

Remark (Outline of the Topic).

> Following the standard pedagogical order:
>
> - **Without uncertainty:** the *Permanent Income Hypothesis* (next section).
> - **With uncertainty:** Hall's *Random Walk Hypothesis*, which adds Euler-equation testable implications (the section after that).
> - **Modern empirical literature:** natural-experiment-based tests of MPC, Excess Sensitivity, excess smoothness, precautionary saving (final section).

1.2 The Permanent Income Hypothesis: No Uncertainty

We begin with the deterministic case. Removing uncertainty is not because it is realistic—it is not—but because it is the cleanest setting in which to derive the central PIH theorem. Uncertainty will be layered on top in the next section.

1.2.1 Setup

The key idea, in one sentence:

utility maximization + a perfect (state-contingent) financial market ⇒ consumption depends only on the net present value of lifetime wealth.

Removing uncertainty from the problem makes the "state-contingent" qualifier vacuous—there is only one state—so what remains is just frictionless intertemporal borrowing and lending at a fixed interest rate.

Simplifying Assumptions

- **Finite horizon:** $t = 0, 1, 2, \ldots, T$.

- **Representative household** with a known, exogenous income process $\{Y_t\}_{t=0}^T$ (no labor-supply choice; income is just an endowment stream).
- **Borrowing and lending** at a constant interest rate r, with full commitment (the household always pays back).
- **Initial assets** A_0 given.

The Household's Problem

The household chooses $\{C_t, A_{t+1}\}_{t=0}^T$ to maximize

$$\boxed{\max_{\{C_t,\, A_{t+1}\}_{t=0}^T} \sum_{t=0}^{T} \beta^t \, u(C_t)}$$

subject to the period budget constraint

$$\frac{A_{t+1}}{1+r} + C_t \;=\; Y_t + A_t, \qquad \forall\, t = 0, 1, \ldots, T,$$

non-negativity $C_t \geq 0$, A_0 given, and the terminal condition

$$A_{T+1} \geq 0 \qquad \text{(natural borrowing limit).}$$

Remark (Reading the Period Budget).

> The form $\frac{A_{t+1}}{1+r} + C_t = Y_t + A_t$ uses the convention that A_t is the household's wealth at the *start* of period t, after last period's interest has accrued. The household then earns income Y_t, consumes C_t, and the remainder is saved at end-of-period in an account that returns $(1 + r)$ next period. Equivalently: the saving made today is worth A_{t+1} tomorrow, so its date-t value is $A_{t+1}/(1+r)$. The terminal condition $A_{T+1} \geq 0$ rules out dying in debt; together with the household's strict preference for consumption ($u' > 0$), it will bind with equality at the optimum—the household leaves no money on the table.

Remark ("Source of the Trade-off").

> Why is there a saving decision at all? Because resources moved across time pay a price/reward: a dollar saved today returns $(1 + r)$ tomorrow, and a dollar borrowed today must be paid back as $(1 + r)$ tomorrow. Marginal utility tomorrow is discounted by $\beta < 1$. The household balances these two forces—the **intertemporal trade-off**—which is exactly what the Euler equation will pin down.

1.2.2 From Period Budgets to the Intertemporal Budget Constraint

The $T+1$ period constraints can be collapsed into a single **intertemporal budget constraint** (IBC) by repeated substitution. Rearrange the period-t constraint:

$$\frac{A_{t+1}}{1+r} - A_t \;=\; Y_t - C_t.$$

Divide by $(1+r)^t$:

$$\frac{A_{t+1}}{(1+r)^{t+1}} - \frac{A_t}{(1+r)^t} \;=\; \frac{Y_t - C_t}{(1+r)^t}.$$

Sum from $t=0$ to $t=T$. The LHS telescopes:

$$\frac{A_{T+1}}{(1+r)^{T+1}} - A_0 \;=\; \sum_{t=0}^{T} \frac{Y_t - C_t}{(1+r)^t}.$$

Rearrange and impose $A_{T+1} \geq 0$:

$$\boxed{\underbrace{\sum_{t=0}^{T} \frac{C_t}{(1+r)^t}}_{\text{PV of lifetime consumption}} \quad \leq \quad \underbrace{A_0 + \sum_{t=0}^{T} \frac{Y_t}{(1+r)^t}}_{\text{PV of lifetime wealth (“Permanent Income” PI)}} \quad .} \tag{IBC}$$

At the optimum (since utility is strictly increasing in C_t), the IBC binds with equality.

Definition 1.1: Permanent Income

The **Permanent Income** of a household is the present value of its total lifetime resources:

$$\text{PI} \;\equiv\; A_0 \;+\; \sum_{t=0}^{T} \frac{Y_t}{(1+r)^t}.$$

It is a single scalar that summarizes everything the household knows about its budget—initial wealth plus the discounted sum of all current and future labor income.

Remark (The IBC Collapses Time).

> The move from $T+1$ period constraints to a single IBC is more than a notational convenience. It says: *from a budgeting standpoint, the household does not face a sequence of separate problems, one per period—it faces a single lifetime allocation problem.* The household can shift consumption arbitrarily across dates as long as the present-value sum is respected. This is what the perfect-financial-market assumption buys you: complete flexibility to move resources through time.

1.2.3 Characterizing the Solution: The PIH Theorem

Maximizing a strictly concave objective subject to a linear constraint gives a clean characterization.

Proposition (Necessary and Sufficient Conditions)

A consumption sequence $\{C_t\}_{t=0}^T$ solves the household's problem if and only if it satisfies:

(1) The **intertemporal budget constraint** with equality:

$$\sum_{t=0}^{T} \frac{C_t}{(1+r)^t} = \text{PI}.$$

(2) The **Euler equation** at every $t = 0, 1, \ldots, T-1$:

$$u'(C_t) = \beta\,(1+r)\,u'(C_{t+1}). \tag{EE}$$

The Euler equation is the standard intertemporal optimality condition: at the optimum, the marginal utility of a dollar consumed today equals the discounted marginal utility of that same dollar saved and consumed tomorrow.

The PIH Theorem

The IBC + EE system has a striking implication.

The Permanent Income Hypothesis

The optimal consumption path $\{C_t\}_{t=0}^T$ depends on the income stream $\{Y_t\}_{t=0}^T$ **only through Permanent Income**.

Why? Because the entire income stream enters the system only via the IBC, and only via the scalar PI. The Euler equation does not involve Y_t at all. So two different income streams that share the same PI lead to the *identical* consumption path—even though their year-by-year profiles can be wildly different.

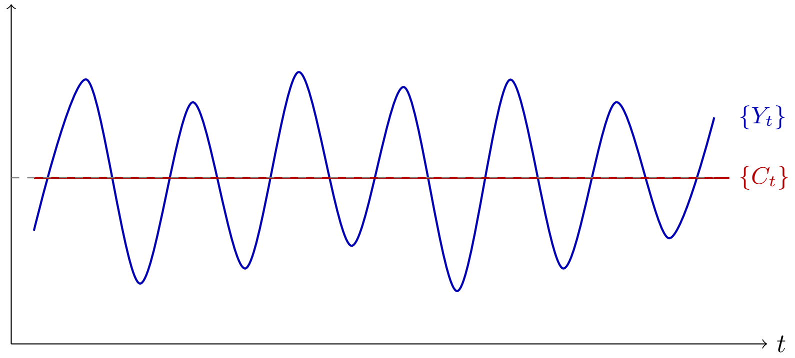

(consumption is constant if $\beta(1+r) = 1$)

The picture is the visual content of PIH: a wildly fluctuating income process can coexist with a perfectly smooth consumption process, because the household uses the saving/borrowing margin to absorb the income fluctuations.

1.2.4 Two Lessons of the PIH

The PIH theorem yields two policy-relevant implications, both deceptively simple.

Fact 1.2: The Timing of Income is Irrelevant

Two income streams $\{Y_t\}$ and $\{Y_t'\}$ that share the same present value PI generate the same consumption path $\{C_t\}$. Receiving income earlier vs. later, in lump sum vs. smoothly, in a recession vs. expansion—none of it matters for consumption, as long as PI is held fixed.

Fact 1.3: Ricardian Equivalence

Suppose the government finances a transfer (subsidy) of g dollars to households today, paid for by a tax of $(1+r)^{\Delta T}\, g$ dollars in ΔT periods. Then PI is unchanged (the transfer's present value equals the tax's present value), so consumption is unchanged. The MPC out of such a fiscal transfer is exactly zero.

Remark (Why "Ricardian"?).

> The principle was articulated (and rejected) by David Ricardo in 1820 and revived by Robert Barro in 1974. In a PIH world, debt-financed government spending is paid for by future taxes; rational households see this and save the entire transfer to pay the future tax. They are, in effect, internalizing the government's budget constraint into their own.
>
> 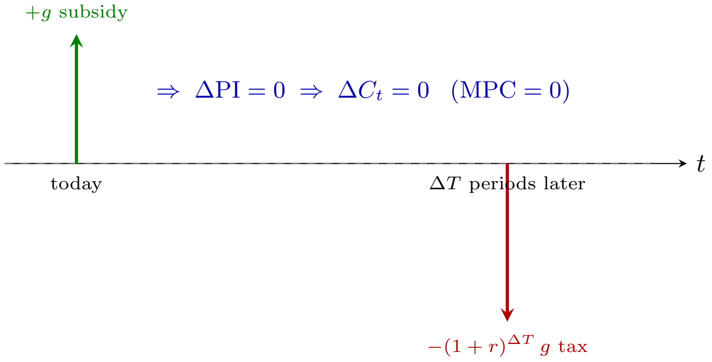
>
>
> Ricardian Equivalence is one of the strongest predictions in macroeconomics, and its rejection in the data is one of the most important pieces of evidence *against* PIH in its purest form. Barro's (1974) revival made it operational by adding bequest motives, which extends the argument to overlapping generations. Empirically, MPCs out of transfers are typically estimated in the range 0.2–0.5, not 0, but the gap shrinks once liquidity constraints, finite lives, and uncertainty are layered on top of PIH.

1.2.5 Example 1: $\beta(1+r) = 1$ and Constant Consumption

The cleanest case parameterizes preferences and the interest rate so that the household weighs the present and the future symmetrically.

Setup

Assume $\beta(1+r)=1$, equivalently $r=(1-\beta)/\beta$.

The Euler equation collapses:

$$u'(C_t)=\beta(1+r)\,u'(C_{t+1})=u'(C_{t+1}) \implies C_t=C_{t+1} \qquad \forall t.$$

Consumption is **constant** across all periods. Plug $C_t=\bar{C}$ into the IBC:

$$\bar{C}\cdot\sum_{t=0}^{T}\frac{1}{(1+r)^t}=\text{PI} \implies \bar{C}=\phi(r)\cdot\text{PI},$$

where the *annuity factor* is

$$\boxed{\phi(r) \;=\; \frac{1-\frac{1}{1+r}}{1-\left(\frac{1}{1+r}\right)^{T+1}}.}$$

The household consumes the annuity value of its lifetime wealth in every period.

Remark (Limits of the Annuity Factor).

> Two limits are worth noting:
>
> - **Infinite horizon ($T\to\infty$):** $\phi(r)\to 1-\frac{1}{1+r}=\frac{r}{1+r}$. The household consumes (roughly) the interest on its wealth and never depletes the principal.
> - **Zero interest ($r\to 0$, finite T):** $\phi(r)\to\frac{1}{T+1}$. The household equally divides total wealth across the $T+1$ periods of life. (This is the Modigliani–Brumberg "life-cycle" intuition in its purest form.)

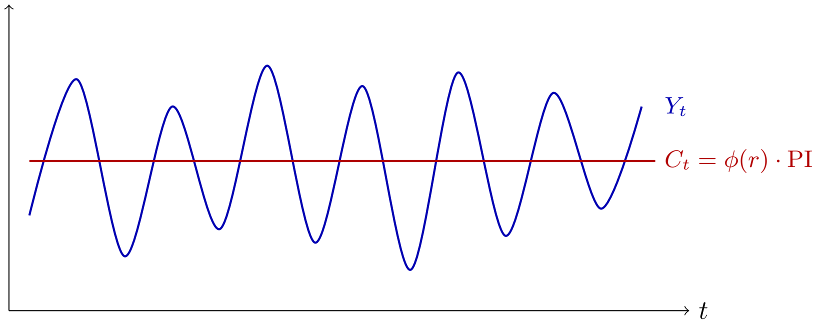

1.2.6 Example 2: General $\beta(1+r)$ and CRRA Preferences

When $\beta(1+r)\neq 1$, consumption is no longer constant. To get a clean closed-form, parameterize utility as constant relative risk aversion (CRRA):

$$u(C)=\begin{cases}\dfrac{C^{1-\sigma}}{1-\sigma} & \text{if } \sigma>0,\ \sigma\neq 1,\\ \ln C & \text{if } \sigma=1,\end{cases} \qquad \sigma>0.$$

Then $u'(C) = C^{-\sigma}$. Plug into the Euler equation:

$$C_t^{-\sigma} = \beta(1+r)\, C_{t+1}^{-\sigma} \iff \left(\frac{C_{t+1}}{C_t}\right)^{\sigma} = \beta(1+r).$$

Take logs:

$$\boxed{g_C \;\equiv\; \ln\!\left(\frac{C_{t+1}}{C_t}\right) \;=\; \frac{1}{\sigma}\,\ln\!\big[\beta(1+r)\big].}$$

Consumption growth is constant and equal to $\frac{1}{\sigma}\ln[\beta(1+r)]$, regardless of t.

Definition 1.4: Intertemporal Elasticity of Substitution (IES)

The parameter $\frac{1}{\sigma}$ is the **intertemporal elasticity of substitution**: the percentage change in the consumption ratio C_{t+1}/C_t per percent change in the gross return $(1+r)$. CRRA preferences impose a constant IES (and a constant coefficient of relative risk aversion equal to σ); these two roles of σ are conflated under CRRA, which **?** showed how to disentangle.

Fact 1.5: Lesson: Consumption Growth Reflects Preferences, not Income

The growth rate of consumption depends on:

- the patience of the household (β),
- the return on saving (r),
- the willingness to substitute consumption across time ($1/\sigma$).

It does **not** depend on the growth rate of income. Two households with the same preferences and same r will display the same consumption growth, even if one's income is rapidly rising and the other's is rapidly falling.

Remark (Three Cases of $\beta(1+r)$).

- $\beta(1+r) = 1$: $g_C = 0$, constant consumption (Example 1).
- $\beta(1+r) > 1$: market rewards saving more than the household discounts the future $\Rightarrow$ $g_C > 0$, consumption rises over time (the household is willing to start poor and end rich).
- $\beta(1+r) < 1$: the household discounts more than the market rewards $\Rightarrow$ $g_C < 0$, consumption falls over time (the household front-loads consumption).

Empirically, U.S. aggregate consumption grows at roughly 2% per year, which combined with $r \approx 0.04$ requires β to be near (but slightly below) 1.

1.2.7 Connecting PIH to the Production Economy

The PIH framework uses an abstract asset A_t with a constant return r. In the production economies of Chapters 5–10, households save in a different form (capital k_t) which is rented out to a firm at rate r_t^k and depreciates at rate δ. These two pictures are equivalent up to a relabeling.

The household's period budget in the production setting is:

$$I_t + C_t \;=\; w_t L_t + r_t^k k_t,$$

where $I_t = k_{t+1} - (1-\delta)k_t$ is gross investment. Substituting:

$$k_{t+1} - (1-\delta)k_t + C_t \;=\; w_t L_t + r_t^k k_t,$$

i.e.,

$$\boxed{k_{t+1} + C_t \;=\; w_t L_t + \tilde{r}_t\, k_t, \qquad \tilde{r}_t \equiv 1 + r_t^k - \delta.}$$

With the identification $A_t \leftrightarrow k_t$, $Y_t \leftrightarrow w_t L_t$, and $1 + r \leftrightarrow \tilde{r}_t$ (the gross return on capital, net of depreciation), this is exactly the asset-economy budget we have been using. Everything we have derived—IBC, Euler equation, PIH, Examples 1 and 2—applies verbatim to the production economy.

Remark (Why the Detour Through Production Matters).

> Two reasons. First, the PIH framework as stated here treats r as exogenous, but in any general-equilibrium model with capital, r is endogenously determined by the marginal product of capital. The IBC is then a single household's budget; in equilibrium, r adjusts so that aggregate consumption and investment add up to aggregate output. Second, this connects PIH to RBC and Aiyagari: those models are GE versions of PIH plus production, plus (in Aiyagari's case) idiosyncratic income shocks. PIH provides the partial-equilibrium consumption function that the GE machinery then aggregates and equilibrates.

1.2.8 Empirical Application: Keynes (1936) and Friedman's Resolution

The PIH was not formulated in a vacuum. **?** wrote it down explicitly to resolve a puzzle posed by **?** and the empirical literature of the next two decades. Walking through the history in some detail does double duty: it shows the kind of empirical reasoning that the PIH framework enables, and it illustrates the broader methodological point that *looking at data through the lens of an economic model often changes what the data appear to say.*

?

In *The General Theory*, Keynes wrote:

> "Consumption mainly depends on current income, and the relation is *fairly stable.*"

This is the so-called **Keynesian consumption function**: $C_i = a + bY_i$, with the marginal propensity to consume (MPC) b a structural parameter. Empirical support came from **cross-sectional** household surveys: in any given year, the projection of consumption on income across households produced a positive, statistically powerful, and seemingly stable slope $\hat{b}$ in the range 0.6–0.8. Richer households did consume more; poorer households consumed less; the relationship looked tight.

For roughly thirty years this was the dominant view of consumption behavior. The MPC of $\hat{b} \approx 0.75$ became a building block in textbook Keynesian multipliers, the IS-LM model, and policy analysis of fiscal stimulus.

The Puzzle

Two empirical facts began to undermine the Keynesian view by the 1940s and 1950s:

- In *long-run time series*, the aggregate saving rate S/Y was approximately *constant* at around 10% across the early 20th century, despite real per-capita income having grown several-fold. A fixed MPC of 0.75 would have predicted an ever-rising saving rate.
- In *short-run time series* (year-to-year aggregate fluctuations), consumption tracked income closely—but with a much smaller MPC than the cross-section implied.

The cross-section, the long run, and the short run gave *three different* answers for the same parameter b. The Keynesian theory, with a single structural MPC, could not accommodate all three.

Friedman's Resolution: Permanent vs. Transitory Income

Friedman's PIH provides the framework. Take Example 1 (Section 7) literally: each household i consumes a constant share of its lifetime wealth,

$$C_{it} = \phi(r)\,\mathrm{PI}_i, \qquad \forall t.$$

Now decompose observed income into a **permanent** and a **transitory** component:

$$Y_{it} \;=\; \underbrace{\phi(r)\,\mathrm{PI}_i}_{\equiv Y_{it}^P \text{ (permanent)}} \;+\; \underbrace{Y_{it}^T}_{\text{(transitory)}} \;.$$

Note that under PIH, C_{it} *is identical* to the permanent component: $C_{it} = Y_{it}^P$. The transitory component Y_{it}^T is what is left over (a windfall, a bonus, a temporary unemployment spell, an unexpectedly good harvest).

Computing Keynes's Coefficient Suppose we run the cross-sectional regression of C_i on Y_i at a fixed time t. The OLS slope is, by construction,

$$\hat{b} \;=\; \frac{\mathrm{Cov}[(,Y]_i\,,C_i)}{\mathrm{Var}\,[(]\,Y_i)}.$$

Substitute $Y_i = Y_i^P + Y_i^T$ and use bilinearity of covariance:

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

Now invoke two structural assumptions implied by the PIH:

- Since $C_i = Y_i^P$, we have $\mathrm{Cov}[(,Y]_i^P,C_i) = \mathrm{Var}[(]\,Y_i^P)$.
- Assume **transitory income is independent of permanent income**: $Y_i^T \perp \mathrm{PI}_i$. Then $\mathrm{Cov}[(,Y]_i^T,C_i) = \mathrm{Cov}[(,Y]_i^T,Y_i^P) = 0$ and $\mathrm{Var}[(]\,Y_i^P + Y_i^T) = \mathrm{Var}[(]\,Y_i^P) + \mathrm{Var}[(]\,Y_i^T)$.

The Keynesian regression coefficient simplifies to

$$\boxed{\hat{b} \;=\; \frac{\mathrm{Var}[(]\,Y_i^P)}{\mathrm{Var}[(]\,Y_i^P) + \mathrm{Var}[(]\,Y_i^T)} \;=\; \text{share of permanent income in cross-sectional income variation.}} \tag{Friedman}$$

Fact 1.6: $\hat{b}$ Is Not the MPC

The coefficient $\hat{b}$ that Keynes interpreted as the MPC is in fact a **signal-to-total ratio**: the share of cross-sectional income variation that comes from permanent differences across households. Under the PIH, the true MPC out of a transitory shock is essentially zero, and the MPC out of a permanent shock is $\phi(r) \approx r$ (small). Neither equals $\hat{b}$.

Why This Resolves the Puzzle

The Friedman decomposition explains all three empirical observations:

- **Cross-section:** Most of the variance *between households* at a point in time is permanent (one household is a doctor, another a janitor; this difference is not transitory). So $\mathrm{Var}[(]\,Y^P)$ is large relative to $\mathrm{Var}[(]\,Y^T)$, $\hat{b}$ is close to 1, and the cross-section delivers a large estimated slope. *Not because households consume out of current income, but because most of the income variation is permanent and fully passes through to consumption.*
- **Long-run time series:** The growth of aggregate income across decades is also permanent. A permanent rise in income of $X\%$ raises consumption by $X\%$, leaving S/Y constant. Consistent with the data.
- **Short-run time series:** Year-to-year aggregate fluctuations have a large transitory component (recessions, booms). PIH predicts that consumption barely responds to these—i.e., a small short-run MPC. Also consistent with the data.

Remark (Visualizing the Decomposition).

> 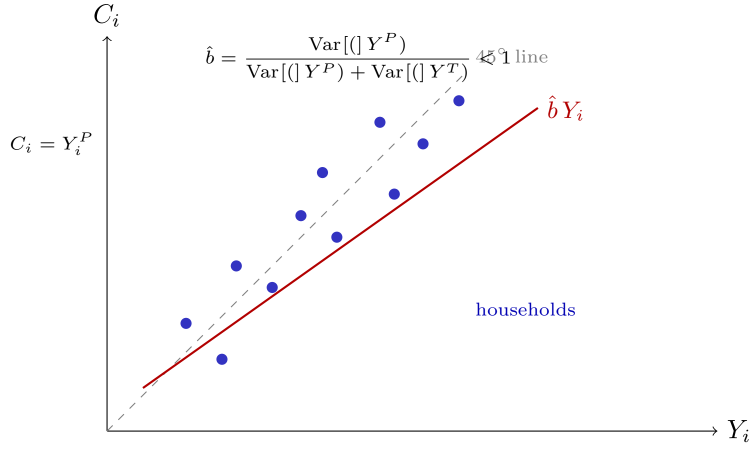
>
>
> Each blue dot is a household. The vertical coordinate $C_i = Y_i^P$ is the permanent part of income; the horizontal coordinate $Y_i = Y_i^P + Y_i^T$ adds transitory noise. The OLS line through the cloud (red) has slope strictly less than 1 because the horizontal noise inflates the denominator $\mathrm{Var}[(]\,Y_i)$ without inflating the numerator $\mathrm{Cov}[(,Y]_i\,, C_i)$. The slope is exactly the signal-to-total ratio.

Remark (A Modern Coda: Measuring Permanent Income is Hard).

> Identifying Y^P and Y^T in real data turns out to be one of the deepest empirical challenges in macro labor. With panel data on income (e.g., the PSID), one can decompose a worker's income time series into transitory and persistent components using statistical assumptions on the income process—but quantifying *subjective* permanent income (what the household believes about its lifetime wealth, taking into account future uncertainty) requires modeling the household's information set and forecasting behavior. Ludvig Straub (Clark Medal 2024) built much of his career on this kind of measurement, including a job-market paper that proposed a way to identify permanent-income shocks from observed updates in saving behavior. Three decades after Friedman, the issue is not closed.

1.3 Adding Uncertainty: Hall's Random Walk Hypothesis

The deterministic PIH is beautifully clean but obviously oversimplified: real income streams are stochastic, and a worker rarely knows even next year's income with certainty. We now make the income process **stochastic** and ask how the household optimizes when it must *forecast* its future. The headline, due to **?**, is that under appropriate assumptions **consumption follows a random walk**: the best predictor of next period's consumption, conditional on everything known today, is today's consumption. This delivers a clean, testable empirical implication that organized two decades of subsequent work.

1.3.1 From PIH to Stochastic Income

What changes when we add uncertainty? Two things, both fundamental.

- **Forecasting becomes part of the problem.** The household no longer knows Y_{t+1}, only its conditional distribution given everything observed up to date t. The Euler equation acquires an expectation operator.
- **Precautionary saving emerges.** If u' is convex (i.e., $u''' > 0$), Jensen's inequality implies $\mathbb{E}_t[u'(C_{t+1})] > u'(\mathbb{E}_t[C_{t+1}])$, so a household facing risky future income saves more than one facing the deterministic-equivalent income. This is invisible in the deterministic PIH and shows up the moment uncertainty is added.

? is essentially the analog of PIH under uncertainty. Its central claim is simple: *in a rational, forward-looking household, the part of consumption growth that can be predicted from time-t information has been smoothed out already.* The only source of consumption changes is unexpected news.

1.3.2 Setup with Stochastic Income

The setup mirrors the deterministic case above but with one substantive change: the income process is now random.

- **Time:** infinite horizon, $t = 0, 1, 2, \dots$
- **Income process:** an exogenous stochastic sequence $\{Y_t\}_{t=0}^{\infty}$. We do not yet specify the law of motion.
- **History:** let

  $$s^t \;\equiv\; \{Y_\tau\}_{\tau=0}^{t}$$

  denote the history of income realizations through date t. The household's information set at t is s^t.
- **Conditional expectation:** we write $\mathbb{E}_t[\cdot]$ as shorthand for $\mathbb{E}[\,\cdot \mid s^t\,]$, expectation conditional on the history known at t.
- **Asset:** a single risk-free bond, with constant gross return $1 + r$. Crucially: there are **no state-contingent claims**.

All choices the household makes at date t are functions of the history s^t, since that is all the household knows. We therefore write consumption and asset holdings as

$$C_t(s^t), \qquad A_{t+1}(s^t).$$

The Household's Problem

Maximize expected discounted lifetime utility:

$$\boxed{\max_{\{C_t(s^t),\, A_{t+1}(s^t)\}} \mathbb{E}_0\left[\sum_{t=0}^{\infty} \beta^t\, u\big(C_t(s^t)\big)\right]}$$

subject to the period budget constraint, which must hold *for every history*:

$$C_t(s^t) + \frac{A_{t+1}(s^t)}{1+r} \;=\; Y_t(s^t) + A_t(s^{t-1}), \qquad \forall\, t,\, \forall\, s^t,$$

non-negativity $C_t(s^t) \geq 0$, A_0 given, and the **no-Ponzi condition**:

$$\lim_{T\to\infty} \frac{A_{T+1}(s^T)}{(1+r)^T} \;\geq\; 0, \qquad \forall\, \{s^T\}_{T\to\infty}.$$

Remark (Why "$A_{t+1}(s^t)$" and not "$A_{t+1}(s^{t+1})$"?).

> This subscript matters more than it looks. Writing A_{t+1} as a function of s^t encodes the assumption that the household must commit to its asset position *before* seeing tomorrow's income realization. Only one bond is available; its return is the same across all states tomorrow, and the household chooses a single quantity at t.
>
> If we had instead allowed $A_{t+1}(s^{t+1})$—i.e., the asset position depends on tomorrow's realized state—we would have given the household a complete set of **state-contingent claims** (Arrow-Debreu securities). It could then guarantee any consumption pattern it wanted across all future states, and the problem would collapse back to the deterministic PIH from the previous section. The whole point of "uncertainty" biting is the absence of state-contingent claims; that is what the index s^t on A_{t+1} encodes.
>
> This distinction is also the reason we need an Euler equation with an expectation operator (Section 4 below) rather than a Euler equation that holds state-by-state.

Remark (Reading the No-Ponzi Condition).

> The no-Ponzi-game condition says: the present value of asset holdings cannot grow without bound, along any history. If it could, the household would be running a Ponzi scheme—borrowing forever and rolling debt at $1+r$, paying it off by borrowing more—and would always strictly prefer to do that, since extra borrowing relaxes today's budget. Ruling this out is necessary for the problem to have a well-defined optimum.
>
> In the deterministic finite-horizon PIH, the analog was simply $A_{T+1} \geq 0$. In the infinite-horizon stochastic case, no-Ponzi must hold along *every* possible history.

1.3.3 Optimality Conditions

Maximization with respect to $A_{t+1}(s^t)$, taking expectations over tomorrow's income realization, yields the **stochastic Euler equation**:

$$\boxed{u'\big(C_t\big) \;=\; \beta\,(1+r)\,\mathbb{E}_t\big[\,u'\big(C_{t+1}\big)\,\big] \qquad \forall\, t = 0,1,2,\ldots} \qquad \text{(Stoch. EE)}$$

This is identical to the deterministic Euler equation *except* that $u'(C_{t+1})$ is now replaced by its conditional expectation. The household equates today's marginal utility (a known quantity) to the expected discounted marginal utility tomorrow.

The terminal condition is the **transversality condition**:

$$\boxed{\lim_{T\to\infty} \beta^T\,\mathbb{E}_0\big[\,u'(C_T)\,A_{T+1}\,\big] \;\le\; 0.} \tag{TVC}$$

The TVC says: the discounted expected value of terminal asset holdings, weighted by marginal utility, must vanish (or be non-positive). It rules out the household being unboundedly wealthy in the limit, which would be wasteful from a utility-maximization perspective.

Remark (Sufficiency: Why TVC + Euler Solves the Problem).

> The Euler equation is a necessary first-order condition; it is satisfied by infinitely many candidate sequences $\{C_t(s^t)\}$, including suboptimal ones that, e.g., consume too little and accumulate unbounded assets. The TVC is what selects the unique optimum among them. Geometrically: the Euler equation pins down the *shape* of consumption (its expected growth rate), and the TVC pins down the *level* (no slack at infinity). Together with the budget constraints, they fully characterize the solution.

1.3.4 The Random Walk Hypothesis (General Form)

?'s (**?**) insight is to recognize the Euler equation as a martingale property. Define the **discounted marginal utility**:

$$X_t \;\equiv\; \beta^t\,u'(C_t).$$

Substitute into the stochastic Euler equation:

$$u'(C_t) = \beta\,(1+r)\,\mathbb{E}_t[\,u'(C_{t+1})\,] \iff \beta^t\,u'(C_t) = (1+r)\,\mathbb{E}_t[\,\beta^{t+1}\,u'(C_{t+1})\,].$$

With $\beta(1+r)$ on both sides, this can be rearranged depending on whether $\beta(1+r)=1$. In the simplest case (which we will treat in Hall's specification below), $\beta(1+r)=1$ and the equation reduces to

$$\boxed{X_t \;=\; \mathbb{E}_t[X_{t+1}].}$$

The Random Walk Hypothesis (General Form)

If $\beta(1+r)=1$, the discounted marginal utility $\{X_t\} \equiv \{u'(C_t)\}$ is a **martingale**: its conditional expectation tomorrow equals its current realization. Equivalently,

$$\mathbb{E}_t\big[\,\Delta X_{t+1}\,\big] \;=\; \mathbb{E}_t\big[X_{t+1} - X_t\big] \;=\; 0.$$

Remark (The Economic Meaning: All Predictable Variation Has Been Smoothed).

> The martingale property is the formal statement of a deep economic intuition. A rational, forward-looking household equates expected marginal utilities across periods. So any *predictable* change in marginal utility—anything the household could see coming using time-t information—would represent an opportunity to do better by consuming more today (or tomorrow). At the optimum, no such opportunities are left:

any expected differences in marginal utility have been smoothed out.

What remains in $X_{t+1} - X_t$ is the **innovation**: the unexpected part, the news. The household optimizes *conditional on its information set*; what it could not have foreseen, it could not have insured against.

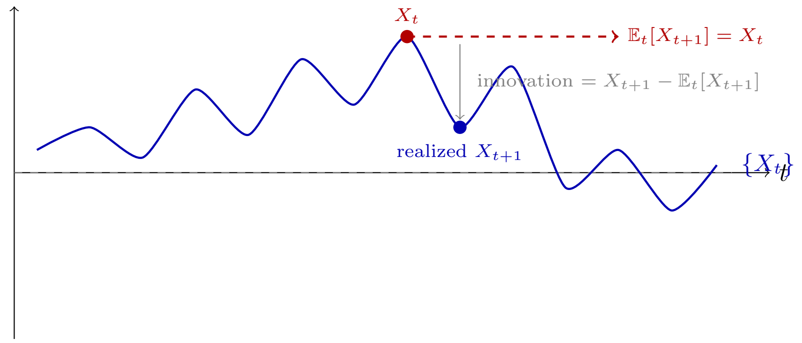

The realized path of X_t wanders (a true martingale), but at every point the conditional expectation of next period's value is exactly X_t itself.

1.3.5 Hall (1978): Quadratic Utility and the Random Walk in Consumption

So far we have shown that *marginal utility* is a martingale—an empirically inconvenient claim, since marginal utility is not directly observable. Hall's contribution is to add two assumptions that make the result testable in terms of the *level of consumption*, which is observable.

Hall's Assumptions

(1) Quadratic utility:

$$u(C) \;=\; C - \frac{a}{2}\,C^2, \qquad a > 0,$$

which gives marginal utility

$$u'(C) \;=\; 1 - a\,C, \qquad u''(C) = -a < 0, \qquad u'''(C) = 0.$$

(2) Patience matches the interest rate: $\beta(1+r) = 1$.

Plug into the stochastic Euler equation:

$$1 - a\,C_t \;=\; \mathbb{E}_t[\,1 - a\,C_{t+1}\,] \;=\; 1 - a\,\mathbb{E}_t[C_{t+1}].$$

The constants and the a cancel:

$$\boxed{C_t \;=\; \mathbb{E}_t[\,C_{t+1}\,].} \qquad \text{(Hall RWH)}$$

?'s (?) Random Walk Hypothesis

Under quadratic utility and $\beta(1+r) = 1$, the level of consumption itself is a martingale: $\{C_t\}_{t=0}^{\infty}$ is a **random walk**. The best predictor of next period's consumption is today's consumption.

Remark (Why Quadratic Utility?).

> The general form of the random walk hypothesis says *marginal utility* is a martingale. To translate this into a statement about *consumption*, we need a relationship between the two:
>
> $$u'(C_t) = \mathbb{E}_t[u'(C_{t+1})] \overset{?}{\Longrightarrow} C_t = \mathbb{E}_t[C_{t+1}].$$
>
> The implication holds if and only if u' is *linear* in C. The unique strictly concave specification with linear u' is the quadratic. Under CRRA, log, or any other "standard" utility, u' is nonlinear, so the unobservable martingale in u' does *not* translate cleanly into a martingale in C. This is the strong assumption Hall is paying for testability.
>
> A consequence of $u'(C) = 1 - a\,C$ is that $u''' = 0$: the third derivative vanishes. By a standard result, this implies **no precautionary saving motive**—the household is risk-averse but not prudent. Future income uncertainty changes the household's consumption distribution but not its level. This is one of the most criticized features of Hall's model and a major reason that later buffer-stock work (**?**) adopted CRRA with $u''' > 0$.

Remark (The Tradeoff: Realism vs. Testability).

> The lecturer's point in introducing Hall's specific assumptions is worth re-emphasizing: when you really want to test whether consumption behavior is rational, what you want to test is whether *marginal utility* is a martingale. That is the structural prediction. Translating this to a test on consumption levels requires a strong utility-function assumption that is itself unlikely to be exactly right. Empirical rejections of Hall's RWH might therefore be rejections of (i) rationality, (ii) $\beta(1+r) = 1$, or (iii) quadratic utility—without further assumptions one cannot disentangle them.

1.3.6 Why It Matters: A Testable Implication

Hall's result is theoretically suggestive, but its real impact came from converting it into a sharp empirical test. If $C_t = \mathbb{E}_t[C_{t+1}]$, then the unexpected change $\varepsilon_{t+1} \equiv C_{t+1} - \mathbb{E}_t[C_{t+1}] = C_{t+1} - C_t$ should be **orthogonal** to all time-t information. In particular, lagged variables X_t in the household's information set should have no predictive power for ΔC_{t+1}:

$$\text{Cov}[(,\Delta]\,C_{t+1},\, X_t) \;=\; 0 \qquad \text{for all } X_t \text{ in the information set.}$$

This is the **Hall test**: regress ΔC_{t+1} on lagged income, lagged consumption growth, lagged stock returns, etc., and look for non-zero coefficients. Under PIH+RWH, all such coefficients should be zero. A statistically significant coefficient indicates **excess sensitivity**—

consumption responding to predictable changes in income—which is the leading robust empirical violation of the model.

Remark (The Modern Empirical Frontier).

> The 1978–1990 literature on Hall's test (**? ?**, Campbell–Mankiw 1989, **? ?**) found systematic excess sensitivity to predictable income changes, particularly for low-wealth households. Modern micro studies using natural experiments—tax rebates, lottery wins, oil-price shocks affecting region-specific incomes—continue to find MPCs out of transitory shocks in the range 0.2–0.5, much higher than the near-zero level predicted by PIH+RWH. The leading explanations involve liquidity constraints (Aiyagari-style buffer-stock saving), present bias (Laibson 1997), and cognitive limitations on forecasting permanent income. Section **??** below walks through the canonical results.

1.3.7 Quadratic Utility and Precautionary Saving

The previous remarks flagged $u''' = 0$ as "the strong assumption Hall is paying for testability." This subsection makes that flag operational: *quadratic utility itself has a testable implication*, namely the **absence of precautionary saving**. The literature has overwhelmingly rejected this implication, which is why modern consumption-saving work uses CRRA preferences and a *drifted* random walk in $\ln C_t$ rather than Hall's clean random walk in levels.

The Mechanism

Return to the stochastic Euler equation, written under $\beta(1+r) = 1$ for clarity:

$$u'(C_t) \;=\; \mathbb{E}_t[\,u'(C_{t+1})\,]\,.$$

Hold $\mathbb{E}_t[C_{t+1}]$ fixed and consider a **mean-preserving spread** of the conditional distribution of $C_{t+1} \mid s^t$ (intuitively, "more uncertainty about tomorrow"). What happens to today's optimal C_t? The answer depends entirely on the curvature of u':

- **If u' is linear in C (i.e., $u''' = 0$, the quadratic case).** Then $\mathbb{E}_t[u'(C_{t+1})] = u'(\mathbb{E}_t[C_{t+1}])$, so the right-hand side of the Euler equation is unchanged by the spread. Today's marginal utility, and therefore today's C_t, is also unchanged. *Pure mean-preserving uncertainty has no effect on optimal consumption today.*

- **If u' is convex in C (i.e., $u''' > 0$, the "prudent" case).** By Jensen's inequality, $\mathbb{E}_t[u'(C_{t+1})] > u'(\mathbb{E}_t[C_{t+1}])$. The mean-preserving spread strictly raises the right-hand side. To restore the equation, $u'(C_t)$ must rise; since $u'' < 0$, this means C_t falls. The household responds to higher uncertainty by **consuming less today and saving more**.

The second case is what we call **precautionary saving**: the household holds extra wealth today specifically to buffer against the variance of future consumption.

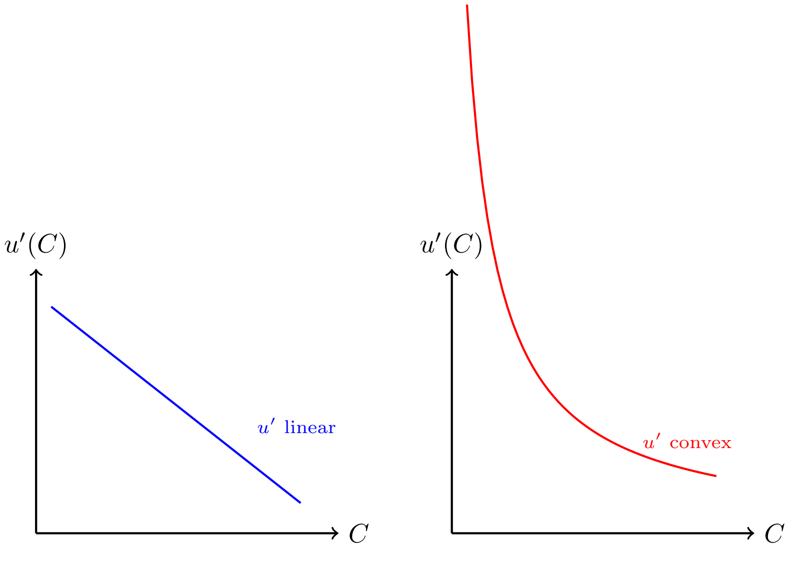

(a) Quadratic utility ($u''' = 0$)
No precautionary saving

(b) Prudent agent ($u''' > 0$)
Precautionary saving present

Remark (Precautionary Saving $\neq$ Risk Aversion).

> A common confusion—surprisingly persistent even among trained economists—is to conflate precautionary saving with risk aversion. They are governed by different derivatives of u:
>
> - **Risk aversion** is about $u'' < 0$: the household dislikes mean-preserving spreads in *utility levels*. A quadratic-utility household is risk-averse—welfare falls when income becomes more variable.
> - **Prudence**, equivalently the precautionary-saving motive, is about $u''' > 0$: the household responds to mean-preserving spreads in *consumption* by *adjusting* how it smooths intertemporally. The Euler equation pins down expected marginal utility, not expected utility, so what governs the saving response is curvature one derivative deeper.
>
> A quadratic-utility household is risk-averse but *not* prudent. It dislikes risk but does not save extra to buffer against it. CRRA, log, and exponential preferences are all both risk-averse and prudent.

Empirical Verdict and a Modern Replacement

The empirical consumption-saving literature has produced *robust* evidence that households increase their saving when their perception of future income risk rises—across many different proxies for risk (job-loss probability, occupation-level income variance, expected health shocks, even children's future labor-market prospects). For one influential example among many, **?** documents what she labels *dynastic precautionary saving*: parents who anticipate higher income risk for their children save more on their own behalf, with implications for overlapping-generations models of redistribution.

Remark (CRRA and the Drifted Random Walk).

> Once $u''' > 0$ is allowed, C_t is no longer a martingale, but a closely related result survives. Under **CRRA preferences** ($u(C) = C^{1-\sigma}/(1-\sigma)$) and a strong assumption on the income-shock distribution (typically lognormal innovations to permanent income), one can derive that *log consumption follows a drifted random walk*:
>
> $$\ln C_{t+1} \;=\; \ln C_t \;+\; \mu \;+\; \nu_{t+1}, \qquad \mathbb{E}_t \nu_{t+1} = 0,$$
>
> where the drift μ depends on β, r, σ, and the variance of innovations. This is the modern formulation that replaces Hall's level random walk in quantitative work, and is the version derived in the homework. The testable implication—that anything in X_t should have zero predictive power for the residual ν_{t+1}—survives the transition; what changes is the functional form of the regression.

1.3.8 The Markov Simplification

The history-based formulation above is conceptually clean but operationally inconvenient: s^t grows without bound. In the chapters on RBC and BCA (and again in the Aiyagari chapter that follows) we used—and will use—the recursive formulation with a fixed-dimensional state instead. The bridge between the two is the **Markov assumption**.

Markov Assumption

Suppose the income process $\{Y_t\}$ is Markov: the conditional distribution of Y_{t+1} depends on s^t only through Y_t. Then for any two histories $\tilde{s}^t, \hat{s}^t$ with the same current realization $\tilde{Y}_t = \hat{Y}_t$ and the same current asset position $\tilde{A}_t = \hat{A}_t$, the household's optimal consumption-saving choice is identical.

In other words, the entire history s^t collapses to two state variables: current income Y_t and current assets A_t. The household's policy can be written as $C_t = g_C(A_t, Y_t)$, $A_{t+1} = g_A(A_t, Y_t)$. This is exactly the recursive structure used in the RBC model (where the state was (a_t, k_t)) and the Aiyagari model (where it was (z_t, k_t) for each household).

Remark (The History Form vs. the Recursive Form).

> The history form $C_t(s^t)$ is more general: it works whether or not income is Markov. The recursive form $C_t = g_C(A_t, Y_t)$ requires the Markov assumption but is dramatically more tractable—it reduces an infinite-dimensional choice problem to a function on a fixed state space. All of the computational machinery (VFI, Aiyagari's outer-loop algorithm, etc.) requires the recursive form.
>
> For empirical work the choice depends on the question: if you only need the Euler equation as a moment condition for testing, the history form is fine and imposes no Markov assumption. If you need to compute the policy function or the wealth distribution, you need Markov plus the recursive formulation.

1.4 Empirical Tests of PIH/RWH

The PIH and the random walk hypothesis are now in hand as theoretical predictions. We turn to the empirical literature that took them seriously as testable claims. The arc of this literature, spanning four decades, can be summarized as follows: aggregate tests rejected the representative-agent rational-expectations model; micro tests using natural experiments tightened the rejection by isolating clean variation in anticipated income; one major paper (Hsieh, 2003) appeared to rescue PIH but was later overturned by a more careful replication (Kueng, 2018). The cumulative verdict is that PIH/RWH systematically understates the consumption response to predictable income changes—the phenomenon now universally called **excess sensitivity**.

1.4.1 The Three Implications to Be Tested

Before reviewing the literature, it helps to organize the predictions into three sharp hypotheses about how consumption C_t should respond to different kinds of income changes ΔY_t:

Fact 1.7: Predictions of PIH/RWH

1. **Anticipated $\Delta Y_t \implies$ no ΔC_t.** If a change in income was already in the household's information set at time $t-1$, it has already been incorporated into Permanent Income, and therefore into C_{t-1}. The realized change today carries no news.

2. **Transitory unanticipated $\Delta Y_t \implies$ small ΔC_t.** An unexpected one-period shock raises Permanent Income only by the annuity value of the shock, which is small. Most of the windfall is saved (and consumed slowly over remaining life).

3. **Permanent unanticipated $\Delta Y_t \implies$ large ΔC_t.** A shock that revises the entire future income path one-for-one (e.g., a persistent productivity shift) raises Permanent Income by approximately the size of the shock itself, and consumption should respond nearly one-for-one.

The strongest and easiest-to-test prediction is (1): under PIH/RWH, consumption should *not* respond to news that was already public at $t-1$. This is the implication around which the entire empirical literature is organized.

In practice no shock is exactly "permanent." The third prediction is operationalized as: persistent unanticipated shocks (e.g., AR(1) innovations with ρ close to 1) generate a consumption response close to the size of the shock.

1.4.2 Aggregate Tests: Hall (1978) and Campbell–Mankiw (1989)

Hall's Original Test

?'s (**?**) original strategy is the obvious one: take aggregate U.S. consumption data, regress its first difference on lagged variables in the household's information set, and check whether anything predicts ΔC_{t+1}. Formally,

$$\Delta C_t \;=\; \alpha \;+\; X_t\,\beta \;+\; \varepsilon_t,$$

where X_t is a vector of variables known to the representative household at time t (e.g., lagged income Y_{t-1}, lagged consumption growth ΔC_{t-1}, lagged stock returns). The null hypothesis is RWH:

$$H_0: \quad \beta = 0.$$

Hall's finding was that for $X_t = (Y_{t-1}, C_{t-1})$, one *statistically* fails to reject $\beta = 0$. He interpreted this as evidence *for* PIH.

Remark (Statistical vs. Economic Significance).

> Failing to reject $\beta = 0$ is not the same as having shown $\beta = 0$. There are two ways the test can come out "negative":
>
> - **(Strong-evidence case.)** The point estimate $\hat{\beta}$ is genuinely close to zero *and* its standard error is small. The confidence interval excludes economically meaningful effects. This is strong evidence for PIH.
> - **(Underpowered case.)** The point estimate is small but the standard error is large. The confidence interval includes economically meaningful effects. This is *not* evidence for PIH; it is just an underpowered test.
>
> Hall's confidence intervals were narrow enough to make a serious argument, but already in his data the upper end of the CI was not literally zero. Subsequent papers focused on this gap.
>
> 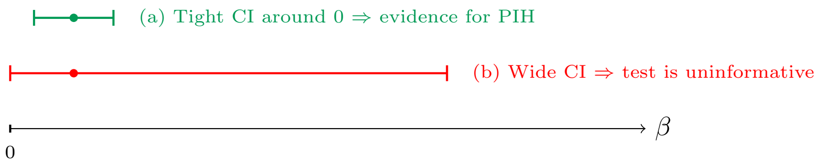
>
>
> Both panels have the same point estimate $\hat{\beta}$, but only (a) constitutes evidence for the null. Hall's and follow-up aggregate work landed somewhere between the two.

?: The Two-Type Model

? sharpened the aggregate test by writing down a structural alternative to PIH. Suppose the population is split into two types:

- A fraction $\lambda \in (0,1)$ of households is **hand-to-mouth**: they consume their income period-by-period, $C_t = Y_t$, so $\Delta C_t = \Delta Y_t$.
- The remaining fraction $1-\lambda$ is fully **PIH**: they smooth, so ΔC_t is pure news, $\Delta C_t = \varepsilon_t$ with $\mathbb{E}_{t-1}\varepsilon_t = 0$.

Aggregating across the two types gives the regression

$$\boxed{\Delta C_t \;=\; \lambda\,\Delta Y_t \;+\; (1-\lambda)\,\varepsilon_t \;+\; X_t\,\gamma, \qquad \mathbb{E}_{t-1}\varepsilon_t = 0.} \tag{C\&M}$$

PIH is the special case $\lambda = 0$. A finding $\hat{\lambda} > 0$ rejects PIH and quantifies the share of consumption that responds to current income—i.e., excess sensitivity.

Remark (The Identification Problem).

> OLS estimation of (C&M) is biased: by construction ΔY_t is correlated with the PIH innovation ε_t (a shock to Y_t *is* the news that updates Permanent Income). Campbell and Mankiw therefore instrument ΔY_t with lagged variables Z_{t-1} that they argue are correlated with ΔY_t but uncorrelated with ε_t—typically lagged stock returns, lagged interest rates, lagged consumption growth. The strategy relies on strong time-series exogeneity assumptions, and the modern econometrics literature views the resulting instruments with considerable skepticism. This is one of the reasons the literature shifted to natural-experiment-based micro tests.

The headline result: $\hat{\lambda} \approx 0.5$ with a standard error of roughly 0.1–0.2. Half of aggregate consumption appears to be hand-to-mouth, decisively rejecting the representative-agent PIH at conventional levels.

Remark (What Aggregate Tests Can and Cannot Show).

> Two caveats are important. First, an aggregate $\hat{\lambda} \approx 0.5$ is consistent with multiple structural stories: a literal 50-50 split between two types, a more complex distribution with most households slightly excessively sensitive, a population with binding liquidity constraints in some periods, a representative agent with present-biased preferences, etc. The aggregate test rejects the *representative-agent rationalization* but does not pin down which structural model takes its place.
>
> Second, the natural correction is to move to micro data, where one can directly observe individual responses to anticipated income changes. The remainder of this section follows that route.

1.4.3 Excess Sensitivity in Micro Data

The micro literature focuses on the strongest PIH prediction: $\partial C_t / \partial(\text{anticipated } \Delta Y_t) = 0$. The empirical strategy is to find natural experiments in which households know about an upcoming income change and to compare consumption before and after the change actually arrives.

A representative sample of the literature:

- **?:** Exploits the U.S. Social Security tax cap, which households hit deterministically partway through the year (after which take-home pay rises mechanically). Compares the consumption of households that hit the cap to those that do not. *Finds non-zero excess sensitivity.*
- **?:** Uses individual variation in the timing of federal tax refunds—refunds are a predictable, anticipated transfer. *Finds non-zero excess sensitivity.*
- **?:** Uses Spanish data on civil servants who receive predictable end-of-year and summer

bonuses. The regular, anticipated structure of the bonus payments is precisely the type of variation PIH says should not move consumption. *Finds non-zero excess sensitivity.*

Across all three studies, the qualitative finding is the same: **ES** $\neq 0$. Households consume measurably more in the period when an anticipated income arrives, even though they could have smoothed perfectly. The micro literature, though heterogeneous in design, was nearly unanimous in this verdict.

1.4.4 Hsieh (2003): Alaska's Permanent Fund

? is the most cleanly identified test in the literature, and for a decade was treated as the definitive evidence *in favor of* PIH. The setup exploits Alaska's Permanent Fund Dividend (PFD), an annuity-style transfer paid to every Alaska resident.

Institutional Background

Alaska's Permanent Fund was created in 1976 as a sovereign-wealth-style fund: 25% of state oil-royalty revenue is deposited into the fund each year. Starting in 1982, 50% of the fund's annual dividend is paid out as the PFD, distributed equally to every individual Alaskan resident regardless of age. The total payout fluctuates year by year with oil-revenue and asset performance, and the per-person amount is determined by dividing the total by the resident count.

Information Timeline—The Source of Identification

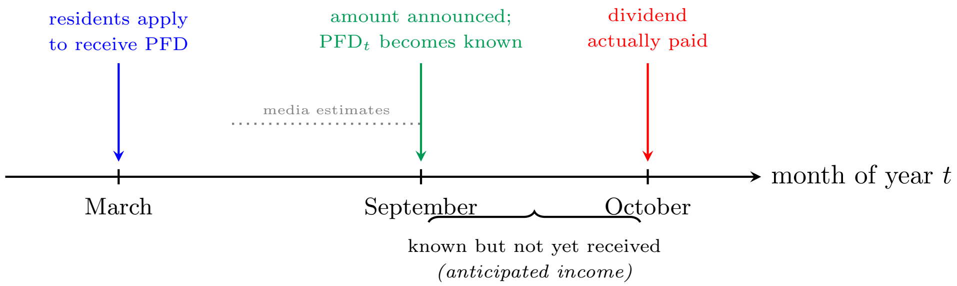

The key timing: by mid-September, every household knows exactly how much PFD it will receive (and rough media estimates have been circulating for months). The actual cash arrives in October. Under PIH, consumption should jump in September (or earlier, when the household first updates its expectations); it should *not* jump in October when an already-anticipated payment is mechanically deposited. The October consumption response is therefore the cleanest possible test of prediction (1).

Specification

For household h in year t, regress the log-change in consumption between Q3 and Q4:

$$\boxed{\ln\left(\frac{C^{Q4}_{t,h}}{C^{Q3}_{t,h}}\right) \approx \alpha \cdot \frac{\text{PFD}_t \times \text{family size}_h}{\text{family income}_{t,h}} + X_{t,h}\,\beta + \varepsilon_{t,h}.} \tag{1.1}$$

The right-hand-side regressor scales the dollar amount of dividend (per-person PFD times family size) by family income, giving the proportional income shock. Three sources of identifying variation feed into this exercise:

- *Across years*: PFD_t varies with oil prices and fund returns.
- *Across households*: family size differs, so the dollar amount differs.
- *Across households*: family income normalizes the shock into a proportional change.

This is effectively a difference-in-differences design with continuous treatment intensity. PIH predicts $\alpha = 0$.

Findings

Spending category	$\hat{\alpha}$	SE
Non-durables	0.0002	0.0324
Durables	0.166	0.088

For non-durables—the category closest to the theoretical C_t—the point estimate is essentially zero, and even three standard errors give an upper bound below 10%. Hsieh interpreted this as a clean failure to reject PIH. For durables, the estimate is meaningfully positive ($\hat{\alpha} \approx 0.17$), but Hsieh argued this is consistent with theory: durables are an investment good, and their bunched purchase in the dividend month is an optimal lumpy response, not a violation of consumption smoothing in the flow-utility sense.

Remark (Why This Was a Big Deal).

> The earlier literature (Parker, Souleles, Browning–Collado) had found excess sensitivity, but each study could be challenged on identification grounds—small samples, weak instruments, idiosyncratic features of the institutional setting. Hsieh's design was unusually clean: large dollar amounts (substantial fraction of household income for many Alaskans), perfect anticipation by October, transparent variation across households, and a sharp prediction. For a decade, the verdict was that the previous excess-sensitivity findings were artifacts and PIH actually held in well-measured data. David Romer's textbook cited Hsieh as state-of-the-art evidence for PIH.

1.4.5 Kueng (2018) Revisits Hsieh

The story did not end there. **?** replicated Hsieh's specification with three modifications, each motivated by a methodological concern:

1. **Normalize by expenditure rather than income.** Income data have well-known measurement problems (under-reporting, missing saving choices, transitory components). Total expenditure is widely viewed as a more reliable proxy for permanent income in micro data, so the proportional shock is rescaled by total household expenditure.

2. **Extend the sample by 12 years.** Hsieh used the original CEX panel covering only the early Alaska Permanent Fund years; Kueng adds more than a decade of additional data, sharply increasing statistical power.

3. **Use non-Alaskans as a control group.** Hsieh's identifying variation came entirely *within* Alaska—across family sizes and incomes. This is potentially problematic if Alaskans have unobserved characteristics (oil-economy seasonality, idiosyncratic Q4 spending patterns) that contaminate the estimate. Kueng adds non-Alaskan households as a control group: people who, by construction, do not receive the PFD and therefore should show no Q3-to-Q4 jump driven by it.

The modification (3) is the most consequential. With non-Alaskans absorbing common Q3-to-Q4 spending patterns and Alaskans absorbing the same patterns plus the PFD response, the difference-in-differences identifies the true PFD effect cleanly.

Findings

After all three modifications, Kueng finds **statistically and economically significant excess sensitivity** for non-durables in Alaska—reversing Hsieh's central conclusion. The PFD, despite being perfectly anticipated by October, generates a measurable consumption response when paid out.

Remark (Identification: A Cautionary Note).

> Kueng's identification is itself imperfect. Comparing Alaskans to non-Alaskans introduces a new concern: any difference between the two groups beyond the PFD could be loaded onto $\hat{\alpha}$. Kueng works hard to control for observable characteristics, but as a matter of principle the tradeoff is between Hsieh's tighter sample (no non-Alaskan contamination, but potentially under-identified) and Kueng's broader sample (cleaner counterfactual, but reliance on observables to align the two groups). This is the natural state of empirical macro—no single design is perfect, and what carries the field forward is whether successive studies converge in spite of their individual identification weaknesses.

1.4.6 Where the Literature Stands

The four-decade arc tells a consistent story once we put all the evidence together:

- Aggregate tests (Hall, Campbell–Mankiw) reject the representative-agent PIH and quantify excess sensitivity at $\hat{\lambda} \approx 0.5$.
- Pre-Hsieh micro tests (Parker, Souleles, Browning–Collado) find non-zero excess sensitivity in cleanly identified natural experiments.
- **?** appeared to overturn this verdict using the Alaska PFD design.
- **?** re-establishes excess sensitivity by adding a control group, extending the sample, and using a more reliable normalization.

The cumulative finding is that anticipated income changes *do* move consumption, in violation of PIH. The leading structural explanations—each motivating a research program in its own right—are:

- **Liquidity constraints.** Households cannot freely borrow against future income. When a transfer arrives, constrained households mechanically increase spending. This is the channel emphasized by buffer-stock-saving models (**?**, **?**) and by Aiyagari-style heterogeneous-agent macro models (next chapter).
- **Present bias.** Hyperbolic discounting (Laibson, 1997) generates excess sensitivity in a representative-agent setting without market frictions.
- **Inattention and information frictions.** Households may not fully process anticipated changes until the cash arrives, blunting forward-looking smoothing.
- **Mental accounting.** Behavioral consumers may treat “windfalls” (tax refunds, dividends) as a separate budget category from regular income.

Remark (The Methodological Lesson).

> Read this less as a settled empirical matter and more as a model of how macro research progresses. Hsieh wrote a beautiful paper that overturned a literature; Kueng wrote an equally beautiful paper that overturned Hsieh; both will likely be overturned in turn. What matters is not that any one estimate is the truth, but that successive iterations of theory and measurement push the field toward a useful approximation. The Aiyagari and HANK frameworks of subsequent chapters can be read as theoretical responses to the empirical regularity, distilled across these papers, that consumption is excessively sensitive to anticipated income.

Remark (Chapter Summary).

> - **The PIH theorem.** With perfect intertemporal financial markets, optimal consumption depends on income only through the present value of lifetime resources (Permanent Income, PI). The timing of income, holding PI fixed, is irrelevant.
> - **Hall’s RWH (1978).** Under quadratic utility and $\beta(1+r)=1$, C_t is a martingale: any predictable change in consumption from time-t information has been smoothed out already. Innovations come only from news about PI.
> - **Quadratic utility is testable, and it fails.** With $u'''=0$ there is no precautionary saving; with $u'''>0$ households cut consumption and save in response to mean-preserving uncertainty—a robust empirical regularity. CRRA + lognormal innovations gives the modern drifted-random-walk formulation.
> - **The empirical literature systematically rejects PIH/RWH.** Aggregate tests yield $\hat{\lambda}\approx 0.5$ hand-to-mouth share; micro tests (Parker, Souleles, Browning–Collado) find non-zero excess sensitivity to anticipated income changes. Hsieh (2003) appeared to rescue PIH using the Alaska PFD, but Kueng (2018) overturned the result with a cleaner control group.
> - **Structural explanations.** Liquidity constraints (Aiyagari-style buffer-stock saving), present bias (Laibson), inattention, and mental accounting all generate excess sensitivity. The Aiyagari and HANK models in the next chapter quantify these channels.

Part III

Problem Sets and Solutions

Subject Index

Author Index