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 9

Business Cycle Accounting

Remark (Notation in This Chapter).

Symbol	Meaning
a_t	efficiency wedge (log-TFP shock, retains the role from Ch. 8)
τ_{Lt}	labor wedge: distortion of the consumption-leisure margin
τ_{It}	investment wedge: distortion of the saving margin
g_t	Government-spending wedge
$\vec{z}_t$	Stacked wedge state vector $(a_t, \tau_{It}, \tau_{Lt}, g_t)$
P, Σ	AR(1) transition matrix and innovation covariance for $\vec{z}_t$
$T(\vec{z}, K)$	Lump-sum rebate of tax revenue
EE_K, EE_L	Distorted inter- and intra-temporal Euler equations

The previous chapter built the canonical RBC model and showed that a single technology shock can qualitatively reproduce the three key business cycle facts. But the qualitative match is far from a quantitative success: at standard parameter values the model under-predicts the cyclicality of hours, mis-times the comovement of consumption and investment, and—most pointedly—requires an implausibly high Frisch elasticity to fit aggregate hours. Real economies clearly violate the frictionless competitive benchmark. The natural next question is: *by how much, and along which margin?*

Business Cycle Accounting (BCA), developed by **?**, henceforth **CKM**, provides a systematic answer. The strategy is to take the RBC model not as a literal description of the economy but as a *measurement device*. We deliberately insert "wedges" into the equilibrium conditions of an otherwise standard RBC model, treat each wedge as an unobserved stochastic process, and let the data tell us which wedges fluctuated when. The resulting decomposition does not say *which* friction caused a recession, but it tells us *which margin* of the frictionless model the recession violated most. That, in turn, sharply narrows the set of structural theories worth taking seriously.

9.1 The Concept: Wedges as Tax-Equivalent Distortions

Suppose we observe in the data that the marginal product of capital is higher than the user cost predicted by the consumption Euler equation. The competitive RBC model predicts these two should be equal. The gap between them is, by definition, a **wedge**: a violation of the optimality condition that we can interpret *as if* an economic agent had been levying a tax on the relevant transaction.

Remark (Wedges are Tax-Equivalents, Not Literal Taxes).

> The wedges in BCA are not interpreted as actual government taxes. They are reduced-form summaries of *any* friction—sticky prices, market power, financial constraints, search costs, information frictions—that drives a wedge between the marginal rate of substitution and the marginal rate of transformation. CKM prove a series of **equivalence theorems** showing that detailed models with very different frictions (e.g., a sticky-price model and a financial-friction model) generate observationally identical paths for $\{C, I, L, Y\}$ to a prototype RBC model with appropriate wedges. BCA exploits this equivalence to back out which margins are distorted, leaving the question of *which* structural friction is responsible for downstream theoretical work.

The Harberger Triangle

The geometry is the standard one from public finance. Consider a single intertemporal margin: the household equates the marginal cost of saving (foregone consumption today, $u'(C)$) to the discounted marginal benefit (extra consumption tomorrow, $\beta u'(C')\,\mathrm{MPK}'$). In the data, however, we observe a gap: at the realized K^{data}, the marginal benefit of one more unit of capital exceeds its marginal cost. The gap is the wedge:

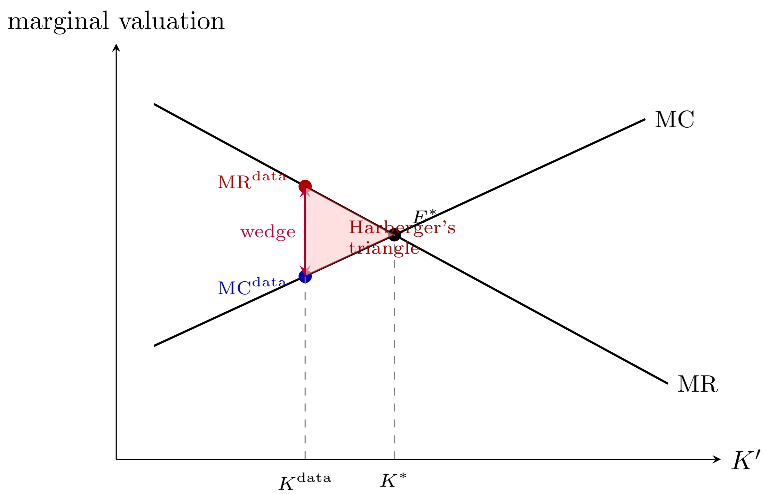

The vertical gap at K^{data} is the marginal wedge—the size of the distortion expressed in the units of marginal valuation. Translated into a tax rate, even a wedge of a few percentage points can represent a substantial tax-equivalent. The shaded triangle is the standard **Harberger triangle**: it measures the deadweight loss generated by the wedge, equal to the area between the MR and MC curves over the distorted region $[K^{\text{data}}, K^*]$.

Remark (Reading the Wedge off an Euler Equation).

In the frictionless RBC model the consumption Euler equation is

$$u'(C_t) = \beta\,\mathbb{E}(u'(C_{t+1})\,\mathrm{MPK}_{t+1})\,.$$

Plug in the data $\{C_t, K_t\}$. Generically, the LHS will not equal the RHS—there will be a residual. The BCA strategy is to attribute the entire residual to a single multiplicative "tax" τ_I on the cost of investment:

$$(1+\tau_{I,t})\,u'(C_t) = \beta\,\mathbb{E}(u'(C_{t+1})\,[\,F_K' + (1+\tau_{I,t+1})(1-\delta)\,])\,.$$

The same logic applies to the intratemporal labor-leisure margin and yields a labor wedge τ_L. The wedges are not estimated structural parameters; they are *whatever* residuals make the equilibrium conditions hold given the observed data. They are by construction reduced-form—which is precisely why they are so informative as a diagnostic device.

9.2 Setup: A Prototype Economy with Stochastic Wedges

The starting point of BCA is a **decentralized RBC model** augmented with three reduced-form distortions:

- **Efficiency wedge** a_t: the standard log-TFP shock from the RBC model. Captures shocks to total factor productivity, but also any unmodeled friction that distorts the production function (e.g., variable capacity utilization, sectoral misallocation).
- **Labor wedge** τ_{Lt}: a tax on labor income that distorts the household's intratemporal labor-leisure choice. Captures any friction that drives a gap between the marginal rate of substitution between consumption and leisure and the marginal product of labor (sticky wages, monopsony, search frictions, distortionary income taxes).
- **Investment wedge** τ_{It}: a tax on the cost of investment that distorts the household's intertemporal saving choice. Captures any friction that drives a gap between the marginal cost and marginal benefit of capital accumulation (financial constraints, collateral requirements, intermediation spreads, capital income taxes).

A fourth process,

- **Government consumption** g_t: stochastic government purchases of goods and services,

is included not so much because government spending is the central object of interest, but because we need a fourth shock for the model to deliver a determinate equilibrium for the four observables $\{C_t, I_t, L_t, Y_t\}$. Without it, the system would be over-identified.

Remark (Why We Need Four Shocks).

> BCA aims to invert the model: given four observed series $\{C, I, L, Y\}$, recover four exogenous processes $\{a, \tau_L, \tau_I, g\}$. For this inversion to be well-posed, the number of shocks must equal the number of observables. Three would leave one observable over-determined; five would leave one shock under-identified. The government-spending wedge g_t is the natural fourth process: it is a real, observable category in the national accounts, but in BCA it functions primarily as the residual that lets the linear system close.

Each of the four processes is assumed to follow an AR(1):

$$z_{t+1}^{(j)} = \rho_j \, z_t^{(j)} + \varepsilon_{t+1}^{(j)}, \qquad \varepsilon^{(j)} \overset{\text{iid}}{\sim} N(0, \sigma_j^2),$$

for $j \in \{a, \tau_L, \tau_I, g\}$. Stacking, the state vector is

$$\boxed{\vec{z}_t \equiv (a_t,\ \tau_{It},\ \tau_{Lt},\ g_t)}$$

with vector AR(1) representation $\vec{z}_{t+1} = P\vec{z}_t + \vec{\varepsilon}_{t+1}$, where P is a 4×4 matrix and $\vec{\varepsilon}_{t+1} \sim N(\vec{0}, \Sigma)$. The matrix P and covariance Σ are the parameters BCA *estimates*; the path of $\vec{z}_t$ is the object BCA *recovers*.

9.3 Recursive Competitive Equilibrium

Because BCA inserts wedges into a *decentralized* environment (rather than working with a planner's problem as in the previous chapter), we must specify the household problem, the firm problem, market clearing, and the government budget separately.

9.3.1 Household

The representative household, taking prices and policy as given, chooses consumption, labor, and capital to maximize lifetime utility. In recursive form:

$$V(\vec{z}, K) = \max_{C,\, L,\, K'} \{u(C, L) + \beta \, \mathbb{E}(V(\vec{z}', K') \mid \vec{z})\}$$

subject to the after-tax budget constraint

$$\boxed{C + (1 + \tau_I)\,[\,K' - (1 - \delta)K\,] \;=\; (1 - \tau_L)\,w(\vec{z}, K)\,L \;+\; r(\vec{z}, K)\,K \;+\; T(\vec{z}, K)}$$

where T is a lump-sum rebate of tax revenue (which we will pin down through the government budget constraint).

The interpretation of each tax:

- $(1 - \tau_L)$ multiplies labor income wL: the household receives only a fraction $(1 - \tau_L)$ of every wage dollar earned.
- $(1 + \tau_I)$ multiplies net investment $I = K' - (1 - \delta)K$: every unit of new capital costs the household $(1 + \tau_I)$ units of consumption goods.

9.3.2 Firm

A competitive representative firm, taking the wage and rental rate as given, hires labor and capital each period:

$$\max_{K,\,L}\; e^{a}\,F(K,L)\;-\;w(\vec{z},K)\,L\;-\;r(\vec{z},K)\,K.$$

The firm operates statically (no capital accumulation decision; that is in the household's problem). The first-order conditions give the standard factor-price equations

$$w(\vec{z},K)=e^{a}F_L(K,L),\qquad r(\vec{z},K)=e^{a}F_K(K,L).$$

9.3.3 Market Clearing and Government Budget

Goods market clearing requires

$$C\;+\;\underbrace{[\,K'-(1-\delta)K\,]}_{=\,I}\;+G\;=\;e^{a}F(K,L).$$

The government budget balances tax revenue against spending plus transfers:

$$G+T=\tau_L\,w(\vec{z},K)\,L\;+\;\tau_I\,[\,K'-(1-\delta)K\,].$$

Remark (Why Lump-Sum Rebates?).

> The lump-sum transfer T is a modeling convenience: it ensures that the wedges τ_L,τ_I act as *pure* distortionary taxes (no income effect from financing government spending). Without it, an increase in τ_L would simultaneously distort labor supply *and* make the household poorer, mixing two channels. With the rebate, raising τ_L only changes relative prices on the margin—a cleaner experiment, and the one CKM use throughout.

9.3.4 Definition

Definition 9.1: Recursive Competitive Equilibrium

A **recursive competitive equilibrium (RCE)** for the prototype economy consists of

- a value function $V(\vec{z}, K)$ and policy functions $g(\vec{z}, K) = (C(\vec{z}, K),\ L(\vec{z}, K),\ K'(\vec{z}, K),\ Y(\vec{z}, K))$,
- pricing functions $w(\vec{z}, K)$ and $r(\vec{z}, K)$,

such that

1. Given prices and policy, the household's problem is solved by V and g;
2. Given prices, the firm's problem is solved at the chosen (K, L);
3. The goods market clears;
4. The government budget balances.

9.4 Wedges in the Equilibrium Conditions

The equilibrium of the prototype economy is conveniently summarized by two distorted Euler equations and the resource constraint. These three equations, together with the AR(1) processes for $\vec{z}_t$, fully characterize the model's predictions for $\{C_t, I_t, L_t, Y_t\}$.

Intratemporal Euler Equation (the labor margin)

The household's first-order condition for labor, combined with the firm's FOC $w = e^a F_L$, yields:

$$\boxed{-u_L(C, L) \;=\; (1 - \tau_L)\, u_C(C, L) \cdot \underbrace{e^a F_L(K, L)}_{\text{MPL}}} \tag{EE_L}$$

Comparing to the frictionless RBC condition $-u_L = u_C \cdot \text{MPL}$, the wedge $(1 - \tau_L)$ enters multiplicatively on the MPL side. Whenever the data show households working "too little" relative to what the wage would justify in a competitive market, BCA registers this as a positive labor wedge $\tau_L > 0$.

Intertemporal Euler Equation (the saving margin)

The household's first-order condition for K' is:

$$\boxed{(1 + \tau_{I,t})\, u_C(C_t, L_t) \;=\; \beta\, \mathbb{E}(u_C(C_{t+1}, L_{t+1})\, [\, e^{a_{t+1}} F_K(K_{t+1}, L_{t+1}) \;+\; (1 + \tau_{I,t+1})(1 - \delta)\,] \,|\, \vec{z}_t)} \tag{EE_K}$$

Two features deserve attention. First, τ_I enters multiplicatively on the cost side ($1 + \tau_{I,t}$ on the LHS) and on the resale value of undepreciated capital tomorrow ($1 + \tau_{I,t+1}$ inside the expectation). Second, the labor wedge τ_L does *not* appear in (EE_K)—only τ_I does.

This separation is what makes the two wedges identifiable from one another: each distorts a different margin.

Remark (Which Margin, Which Wedge?).

A useful mnemonic:

- The **labor wedge** τ_L distorts the *within-period* trade-off between consumption and leisure. It shows up in (EE_L).
- The **investment wedge** τ_I distorts the *across-period* trade-off between consumption today and consumption tomorrow. It shows up in (EE_K).
- The **efficiency wedge** a_t distorts the production function itself. It shows up wherever F does, which is in both Euler equations and the resource constraint.
- The **government wedge** g_t enters only through the resource constraint (a demand component that competes with C and I).

9.5 Estimation: How CKM Recover the Wedges

The model has now been transformed into a state-space system: four observables $\{C_t, I_t, L_t, Y_t\}$, four latent stochastic processes $\vec{z}_t = (a_t, \tau_{It}, \tau_{Lt}, g_t)$, and equilibrium conditions linking them. The estimation strategy is canonical for such systems.

The ? Estimation Procedure

Step 1: Calibrate non-stochastic parameters. Fix $\beta, \delta, \alpha, \sigma, \nu$, and the form of the utility and production functions to standard values from the RBC literature.

Step 2: Take HP-filtered data. Detrend $\{C_t, I_t, L_t, Y_t\}$ and work with their cyclical components.

Step 3: Log-linearize the equilibrium conditions around the deterministic steady state. The resulting system is linear in $\{\hat{C}_t, \hat{I}_t, \hat{L}_t, \hat{Y}_t\}$ and $\{\hat{a}_t, \hat{\tau}_{It}, \hat{\tau}_{Lt}, \hat{g}_t\}$.

Step 4: Estimate P and Σ by maximum likelihood. Treat the AR(1) parameters (P, Σ) as unknowns. Use the Kalman filter to compute the likelihood of the observed data given the model and parameters. Maximize over (P, Σ).

Step 5: Recover the wedge series. With $(\hat{P}, \hat{\Sigma})$ in hand, run the Kalman smoother to back out the most likely path of $\vec{z}_t$ given the entire observed sample. This yields the estimated wedges $\{\hat{a}_t, \hat{\tau}_{Lt}, \hat{\tau}_{It}, \hat{g}_t\}_t$.

Remark (Why Maximum Likelihood and Not OLS?).

Naively, one might try to compute $\hat{\tau}_{Lt}$ at each date by plugging $\{C_t, L_t, K_t\}$ into the intratemporal Euler equation and solving for τ_{Lt}. This works in principle but ignores the cross-equation restrictions: in a rational-expectations equilibrium, the household's choice of C_t depends on its forecast of *all four* future wedges, and that forecast depends on the joint AR(1) parameters (P, Σ). The MLE/Kalman approach uses the full system to extract wedges efficiently and to estimate (P, Σ) jointly. It also gives standard errors and lets us run counterfactuals (e.g., "what would output have done if only the efficiency wedge had moved?").

9.6 What BCA Has Found

The headline result of **?**, and of the large BCA literature that followed, is striking and remarkably robust across countries and time periods:

The labor wedge does most of the work.

More specifically:

- During major U.S. recessions—the Great Depression of 1929–39, the 1982 recession, the Great Recession of 2008–09—the bulk of the cyclical decline in output is attributable to movements in the **efficiency wedge** a_t and the **labor wedge** τ_{Lt}, with the labor wedge typically explaining the larger share of hours and a comparable share of output.
- The **investment wedge** τ_{It} is surprisingly small. Even in episodes that look superficially like "credit crunches" (the Great Depression, 2008), the investment wedge plays a minor role. This was one of the most controversial findings of CKM, since it suggested that financial-friction models—then ascendant in macro—were missing the main story.
- The **government-spending wedge** g_t explains very little of cyclical fluctuations in $\{C, I, L, Y\}$, consistent with the prior view that fiscal shocks are not the dominant driver of postwar U.S. business cycles.

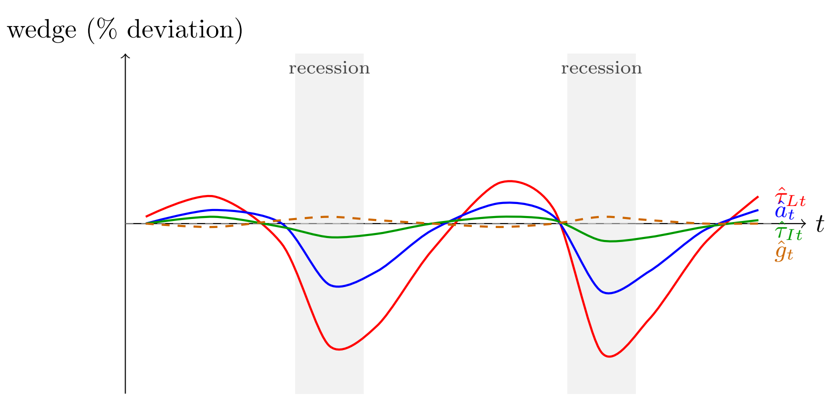

The figure above is a stylized rendering of the CKM result: in each NBER-dated recession (gray bars), the labor wedge falls sharply, the efficiency wedge falls moderately, the investment wedge barely moves, and the government wedge is essentially flat.

Remark (The Ranking Matters More Than the Levels).

The exact percentages CKM report depend on the calibration, the sample, the detrending method, and the specification of the production function. But across many sensitivity checks—and across many countries (**? ?** collected studies of dozens of recessions worldwide)—the qualitative ranking $\hat{\tau}_L \succ \hat{a} \succ \hat{\tau}_I \succ \hat{g}$ in importance is remarkably stable. This is the durable lesson of BCA, even if any single point estimate can be debated.

Implication: Where Theory Should Focus

If the labor wedge is what fluctuates most over the cycle, then the structural friction we are looking for is one that (i) primarily distorts the consumption-leisure margin, and (ii) varies systematically over the cycle. Candidate stories that pass this filter include:

- **Sticky wages.** Nominal wage rigidity, as in New Keynesian models, generates a time-varying gap between MRS and MPL. When aggregate demand falls, real wages remain too high, so employers hire too little and the labor wedge widens.
- **Search and matching frictions.** In Diamond–Mortensen–Pissarides labor markets, wages reflect a Nash bargain between workers and firms rather than the marginal product, and the bargain shifts with labor-market tightness over the cycle.
- **Distortionary taxes and transfers.** Cyclical changes in marginal tax rates, unemployment insurance, or means-tested transfers can mechanically generate a cyclical labor wedge.
- **Monopsony / monopoly markups.** Time-varying market power directly produces a wedge between MPL and the wage paid.

Conversely, models whose primary mechanism distorts the *investment* margin (e.g., pure financial friction models in the spirit of Bernanke–Gertler) face a quantitative challenge: BCA says the investment wedge does not move much, so a friction operating exclusively through that margin will not explain the bulk of cyclical fluctuations. (This does not rule out financial frictions altogether: a financial friction might amplify shocks via labor-market channels, in which case BCA would attribute the action to τ_L, not τ_I.)

9.7 Limitations and Critiques

BCA is a measurement device, not a structural model. Several caveats are important to bear in mind.

- **Wedges are not structural.** BCA tells us *which margin* is distorted, not *what causes* the distortion. Different deep theories can produce identical wedge patterns (this is the equivalence theorem). Moving from a BCA finding to a structural theory still requires additional identifying assumptions.
- **Specification dependence.** **?** showed that BCA results can be sensitive to the assumed period utility function and to whether some wedges are restricted to enter as multiplicative taxes vs. as additive shifters. The qualitative ranking is robust; the precise decomposition is not.
- **Linearization.** Standard BCA log-linearizes around the steady state. In severe recessions or near a binding constraint, the linearization may misattribute nonlinear dynamics to one wedge or another. Recent work has extended BCA with nonlinear filters.

- **The labor wedge is not a clean object.** A growing literature (**? ?**; **? ?**) has shown that decomposing τ_L further—into a household side (the gap between MRS and the wage) and a firm side (the gap between the wage and MPL)—reveals that most of the cyclical labor wedge is on the household side. This has reignited the debate about whether the labor wedge reflects sticky wages, mismeasurement of the consumption-leisure trade-off, or something else entirely.

Remark (BCA in Perspective).

> The deeper contribution of BCA is methodological. It taught macroeconomists to ask, before writing down a detailed structural model: *which margin of the frictionless benchmark is the one I am trying to explain?* A model whose mechanism operates on a margin that the data say is undistorted is, from a quantitative perspective, a non-starter—no matter how elegant. Conversely, any model that successfully matches a margin BCA flags as cyclically distorted is at least in the right neighborhood. This has shaped the agenda of macro for two decades: from sticky-price models that endogenize the labor wedge, to search models that microfound it, to recent work on heterogeneous-agent New Keynesian (HANK) models that connect the labor wedge to inequality and the marginal propensity to consume.

Remark (Chapter Summary).

> - **BCA repurposes the RBC model.** Insert four reduced-form "wedges" (efficiency, labor, investment, government) into the equilibrium conditions and back them out from the data via Kalman filtering of a linearized state-space system.
> - **Wedges are tax-equivalents, not literal taxes.** Equivalence theorems (CKM 2007) show that wedges with appropriate stochastic processes can replicate paths of $\{C, I, L, Y\}$ generated by detailed structural models with very different micro frictions.
> - **The labor wedge does most of the work.** Across decades, countries, and detrending choices, $\hat{\tau}_L$ accounts for the bulk of cyclical variation in major recessions; the investment wedge $\hat{\tau}_I$ is small. The government wedge contributes little.
> - **Implications for theory.** Models whose primary mechanism distorts the consumption-leisure margin (sticky wages, search and matching, monopsony) survive the BCA filter. Pure financial-friction models that operate exclusively through τ_I do not.
> - **Limitations.** Wedges are reduced-form, not structural; results depend on log-linearization; the labor wedge can be further decomposed into household-side and firm-side components, and most of the cyclicality is on the household side.

Part III

Problem Sets and Solutions

Subject Index

Author Index