# ECON 607: Macroeconomic Theory II: Lecture Notes

May 14, 2026

## 1 A Simple Dynamic Exchange Economy

We study a finite-horizon pure exchange economy with time indexed by $t = 0, 1, \dots, T$, where $T$ is odd. There is a single non-storable, non-durable consumption good in every period.

### 1.1 Agents, Endowments, and Preferences

There are two types of agents, indexed by $i \in \{A, B\}$, with populations $N^A$ and $N^B$. Aggregate uncertainty is absent.

Endowments alternate deterministically over time:

- In even periods $t$, type $A$ agents receive one unit of the good and type $B$ agents receive zero.
- In odd periods $t$, type $B$ agents receive one unit of the good and type $A$ agents receive zero.

Preferences are time-separable and identical across agents. Each agent maximizes

$$\sum_{t=0}^{T} \beta^t \ln c_t,$$

where $c_t$ denotes individual consumption in period $t$ and $\beta \in (0, 1)$.

### 1.2 Assets and Budget Constraints

Trade takes place through a one-period nominal bond. Let $P_t$ denote the price of the consumption good in period $t$, $i_t$ the nominal interest rate between $t$ and $t+1$, and $s_t^i$ nominal savings of agent $i$ chosen in period $t$.

We assume agents start with zero assets.

**Period 0.**

$$P_0 y_0^i \geq P_0 c_0^i + s_0^i.$$

**Periods $0 < t < T$.**

$$P_t y_t^i + s_{t-1}^i (1 + i_{t-1}) \geq P_t c_t^i + s_t^i.$$

**Terminal period $T$.**

$$P_T y_T^i + s_{T-1}^i (1 + i_{T-1}) \geq P_T c_T^i, \qquad s_T^i \geq 0.$$

The terminal non-negativity constraint on savings rules out Ponzi schemes.

## 1.3 Real Bonds and the Fisher Equation

Define real bond holdings as

$$b_t^i \equiv \frac{s_t^i}{P_t}.$$

Dividing the budget constraint by $P_t$ and multiplying by $P_{t-1}/P_{t-1}$ yields

$$y_t^i + b_{t-1}^i(1 + r_{t-1}) \geq c_t^i + b_t^i,$$

where the real interest rate satisfies the Fisher equation

$$1 + r_t \equiv \frac{P_t}{P_{t+1}}(1 + i_t) = \frac{1 + i_t}{1 + \pi_{t+1}}.$$

## 1.4 Competitive Equilibrium

**Definition.** A competitive equilibrium consists of a sequence of real interest rates $\{r_t\}_{t=0}^{T-1}$ and allocations $\{c_t^A, c_t^B, b_t^A, b_t^B\}_{t=0}^{T}$ such that:

1. Taking prices as given, each agent maximizes utility subject to their budget constraints.
2. Goods markets clear:

$$N^A y_t^A + N^B y_t^B = N^A c_t^A + N^B c_t^B.$$

3. Bond markets clear:

$$N^A b_t^A + N^B b_t^B = 0.$$

Because the good is non-storable, aggregate saving is zero in equilibrium.

## 1.5 Household Optimization and Euler Equation

Each agent solves

$$\max_{\{c_t, b_t\}} \sum_{t=0}^{T} \beta^t \ln c_t$$

subject to the sequence of real budget constraints and $b_T \geq 0$.

The Lagrangian yields the first-order conditions

$$\frac{\beta^t}{c_t} = \lambda_t, \qquad \lambda_t = \lambda_{t+1}(1 + r_t),$$

which imply the Euler equation

$$\frac{1}{c_t} = \beta(1 + r_t)\frac{1}{c_{t+1}}.$$

In the terminal period, the no-Ponzi condition binds, so $b_T^i = 0$.

## 1.6 Consumption Growth and Risk Sharing

Since all agents face the same interest rate and have identical preferences, their consumption grows at the same rate:

$$\frac{c_{t+1}^A}{c_t^A} = \frac{c_{t+1}^B}{c_t^B}.$$

Therefore, individual consumption growth must equal aggregate consumption growth.

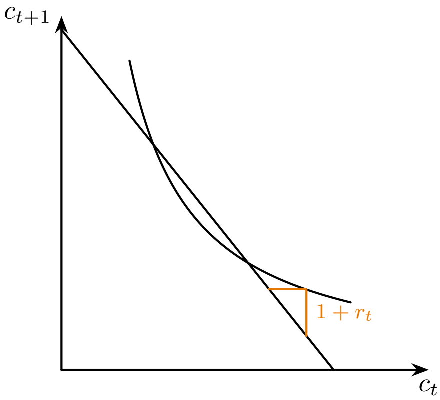

**Proposition.** *In equilibrium, consumption growth equals the growth rate of aggregate endowments.*

**Intuition.** With no storage and complete risk sharing through bond markets, the only feasible intertemporal allocation smooths consumption subject to aggregate resources.

## 1.7 Equilibrium Interest Rates

Aggregate endowments alternate over time:

$$Y_t = \begin{cases} N^A & \text{if } t \text{ is even,} \\ N^B & \text{if } t \text{ is odd.} \end{cases}$$

Thus,

$$\frac{C_{t+1}}{C_t} = \begin{cases} \dfrac{N^B}{N^A} & \text{if } t \text{ is even,} \\ \dfrac{N^A}{N^B} & \text{if } t \text{ is odd.} \end{cases}$$

Using the Euler equation,

$$1 + r_t = \begin{cases} \dfrac{1}{\beta}\dfrac{N^B}{N^A} & \text{if } t \text{ is even,} \\ \dfrac{1}{\beta}\dfrac{N^A}{N^B} & \text{if } t \text{ is odd.} \end{cases}$$

The real interest rate oscillates deterministically around $1/\beta$.

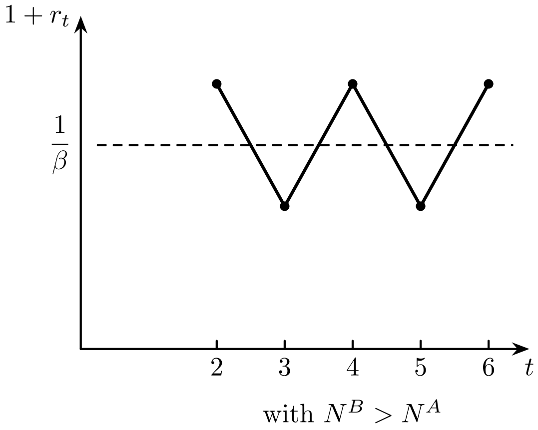

The intuition behind this oscillation is best understood by asking which group is the *net lender* in each period. The Euler equation tells us that the real interest rate—the intertemporal price of consumption—adjusts so as to clear the intertemporal market in every date $t$, i.e., to equate the supply and demand for loanable funds coming from groups $A$ and $B$.

To begin with, the rate is naturally anchored around $1/\beta$. In a frictionless benchmark, the equilibrium intertemporal return is pinned down by impatience: the aggregate desire to shift consumption across time is disciplined by the discount factor $\beta$, so the relevant reference point for the gross return is $1/\beta$.

In the present environment, we depart from that benchmark because (i) the two groups have different sizes, with $N^B > N^A$, and (ii) endowments alternate deterministically over time. As a result, the identity of the net lender alternates across dates, and the real interest rate must move to clear the market for loanable funds.

When group $B$ is the one receiving the endowment (and hence has resources to lend), it becomes the net supplier of loanable funds. Since $B$ is larger than $A$, the supply of savings is high relative to the demand for borrowing by group $A$ (which needs funds to smooth consumption). This relative abundance of loanable funds puts *downward* pressure on the equilibrium interest rate.

Conversely, when group $B$ does *not* receive the endowment, it becomes a net borrower and must obtain funds from group $A$ (the group currently endowed). In that case, demand for loanable funds is high relative to supply, because the large group $B$ is now on the borrowing side while the lending capacity comes from the smaller group $A$. This relative scarcity of loanable funds pushes the equilibrium interest rate *upward*.

In short, the real interest rate oscillates because the economy alternates between periods of *abundant* loanable funds (when the large group is lending) and periods of *scarce* loanable funds (when the large group is borrowing), while the level around which it fluctuates is disciplined by the intertemporal benchmark $1/\beta$.

# 2 Intertemporal Budget Constraint and Extensions

## 2.1 The Intertemporal Budget Constraint

Starting from the period-by-period real budget constraint,

$$c_t + b_t = y_t + (1 + r_{t-1})\, b_{t-1},$$

we can solve for lagged bond holdings as

$$b_{t-1} = \frac{c_t + b_t - y_t}{1 + r_{t-1}}.$$

In particular, under a *no-bequest* (no terminal asset) condition $b_T = 0$, the terminal-period constraint implies

$$b_{T-1} = \frac{c_T - y_T}{1 + r_{T-1}}.$$

Using the constraint in period $t-1$,

$$c_{t-1} + b_{t-1} = y_{t-1} + (1 + r_{t-2})\, b_{t-2},$$

and substituting the expression for $b_{t-1}$ yields

$$c_{t-1} + \frac{c_t + b_t - y_t}{1 + r_{t-1}} = y_{t-1} + (1 + r_{t-2})\, b_{t-2}.$$

Equivalently,

$$c_{t-1} + \frac{c_t}{1+r_{t-1}} = y_{t-1} + \frac{y_t}{1+r_{t-1}} + (1+r_{t-2})\, b_{t-2} - \frac{b_t}{1+r_{t-1}}.$$

Iterating this argument forward from $t = 0$ to $t = T$ delivers the finite-horizon intertemporal budget constraint with a terminal asset term:

$$c_0 + \sum_{t=1}^{T} \frac{c_t}{\prod_{s=0}^{t-1}(1+r_s)} + \frac{b_T}{\prod_{s=0}^{T-1}(1+r_s)} = y_0 + \sum_{t=1}^{T} \frac{y_t}{\prod_{s=0}^{t-1}(1+r_s)} + (1+r_{-1})\, b_{-1}.$$

Imposing the no-bequest condition $b_T = 0$ (and, if desired, $b_{-1} = 0$ for zero initial wealth) yields the simplified version:

$$c_0 + \sum_{t=1}^{T} \frac{c_t}{\prod_{s=0}^{t-1}(1+r_s)} = y_0 + \sum_{t=1}^{T} \frac{y_t}{\prod_{s=0}^{t-1}(1+r_s)}.$$

Thus, with no terminal assets (and no initial wealth), lifetime consumption equals the present discounted value of endowments.

## 2.2 Consumption Levels Across Agent Types

**Remark.** In equilibrium, individual consumption depends on the parity of time and the agent type.

For type $A$,

$$c_t^A = \begin{cases} c_{\text{odd}}^A & \text{if } t \text{ is odd,} \\ c_{\text{even}}^A & \text{if } t \text{ is even,} \end{cases}$$

with an analogous expression for type $B$.

**Intuition.** Endowments alternate deterministically across agent types. With complete markets and identical preferences, agents smooth consumption intertemporally, but consumption levels still reflect systematic differences in income timing. Figure 1 illustrates the alternating endowment process for type $A$ and the resulting smoothed (but still oscillatory) consumption path.

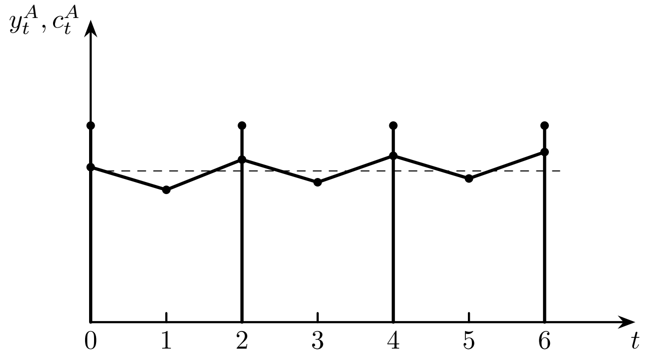

Figure 1: Alternating endowments and smoothed (yet oscillatory) consumption for type $A$.

The key point is that the environment is *two-period deterministic*: endowments switch by parity, and (because $N^B \neq N^A$) the aggregate resources available in the economy also switch by parity. A one-period bond in zero net supply can only *reallocate* resources across agents within a date; it cannot shift aggregate resources across time. Hence, if aggregate endowment is higher in (say) even dates than in odd dates, aggregate consumption must be higher in even dates than in odd dates, and each agent's consumption inherits a two-period cycle.

Formally, let aggregate endowment at date $t$ be

$$Y_t \;\equiv\; N^A y_t^A + N^B y_t^B .$$

Goods market clearing implies, period by period,

$$N^A c_t^A + N^B c_t^B \;=\; Y_t \qquad \forall t.$$

With complete markets and identical preferences, equilibrium risk sharing implies a time-invariant marginal-utility relationship across types; in the symmetric case this collapses to

$$u'(c_t^A) = u'(c_t^B) \quad \Rightarrow \quad c_t^A = c_t^B \equiv c_t,$$

so that

$$c_t \;=\; \frac{Y_t}{N^A + N^B}.$$

Therefore, whenever $Y_t$ alternates deterministically (which occurs here because endowments alternate across types and the population sizes differ), consumption must alternate deterministically as well:

$$c_{\text{even}} \neq c_{\text{odd}} \quad \Rightarrow \quad c_t \text{ oscillates with period 2.}$$

(If instead $Y_t$ were constant over time—e.g., if the two populations had the same size or endowments did not generate aggregate fluctuations—then the complete-markets allocation would deliver constant consumption and there would be no oscillation.)[1]

### 2.3 Extension I: Infinite Horizon

Let $T \to \infty$. The intertemporal budget constraint becomes well-defined only if a no-Ponzi condition is imposed:

$$\lim_{T\to\infty} \frac{b_T}{\prod_{s=0}^{T}(1+r_s)} \geq 0.$$

In the household optimization problem, the corresponding transversality condition is

$$\lim_{T\to\infty} \lambda_T b_T = 0,$$

which rules out explosive accumulation of assets or debt.

**Interpretation (No-Ponzi vs. Transversality).** When $T \to \infty$, the intertemporal budget constraint contains a terminal term (the present discounted value of wealth at $T$). Without further restrictions, an agent could finance any consumption path by rolling over debt forever: the sequence of budget constraints would not impose a meaningful feasibility restriction on lifetime consumption.

The *no-Ponzi* condition rules out precisely this pathology. Written as

$$\lim_{T\to\infty} \frac{b_T}{\prod_{s=0}^{T}(1+r_s)} \geq 0,$$

it says that the present value of terminal bond holdings cannot be arbitrarily negative in the limit. Equivalently, debt cannot grow fast enough (relative to the discounting induced by interest rates) to sustain consumption via perpetual refinancing.

---

[1] All this story is consistent with the growth rates of consumption derived at the begining. Sometimes they are more than 1, sometimes less than 1. That's all.

The *transversality condition* (TVC),

$$\lim_{T\to\infty} \lambda_T b_T = 0,$$

is the optimality counterpart of this idea. Here $\lambda_T$ is the shadow value of wealth (the Lagrange multiplier on the period-$T$ budget constraint), so $\lambda_T b_T$ is the utility value of the agent's terminal portfolio. The TVC rules out two inefficient behaviors: (i) accumulating unbounded assets that are never used for consumption (leaving "money on the table"), and (ii) financing consumption through explosive debt positions that violate intertemporal feasibility.

In standard settings, the no-Ponzi condition is imposed as a feasibility constraint, while the TVC is implied by optimality (together with regularity conditions). Intuitively: no-Ponzi prevents "free lunches" via debt rollovers, and the TVC ensures that the optimal plan does not hoard wealth forever.

## 2.4 Extension II: Storage Technology

Suppose now that storage is available. Endowments are given by

$$y_t^A = \{1, \varepsilon, 1, \varepsilon, \dots\}, \qquad y_t^B = \{0, 0, 0, 0, \dots\},$$

where $0 < \varepsilon < 1$.

Agents can store the good using a storage technology:

$$x_t \longrightarrow x_{t+1}(1-\delta),$$

with depreciation rate $\delta \in (0,1)$ and capacity constraint

$$x_t \leq \bar{x}.$$

Let $\rho_t$ denote the per-unit storage price paid at date $t$.

**Equilibrium Modifications.** The definition of competitive equilibrium must be augmented to include:

- An additional market for storage,
- Storage choices in the budget constraint,
- An additional price sequence $\{\rho_t\}$,
- Modified market-clearing conditions.

The goods market clearing condition becomes

$$N^A y_t^A + N^B x_{t-1}^B(1-\delta) = N^A c_t^A + N^B c_t^B + N^B x_t^B,$$

or equivalently,

$$y_t = c_t + N^B\big(x_t^B - x_{t-1}^B(1-\delta)\big),$$

where the last term represents the change in inventories and plays the role of investment.

**Unlimited Storage.** If $\bar{x} \to \infty$, then in equilibrium:

$$\rho_t \to 0, \qquad b_t \to 0,$$

as storage perfectly smooths consumption over time.

**Interpretation (What changes when storage is available?).** Allowing storage adds a *physical* way of transferring resources across time. In the baseline exchange economy, the one-period bond is in zero net supply and only reallocates goods *across agents within a date*; it cannot smooth *aggregate* resources across time. Storage breaks this restriction: by carrying the good from $t$ to $t+1$ (subject to depreciation), the economy can reallocate aggregate resources intertemporally via inventories.

This is why the goods-market clearing condition acquires an "investment" term. The wedge

$$N^B\big(x_t^B - x_{t-1}^B(1-\delta)\big)$$

is the change in inventories: goods put into storage at date $t$ reduce resources available for contemporaneous consumption, while depreciated goods carried from $t-1$ expand resources available at date $t$. This term plays the same accounting role as investment in a production economy.

**Equilibrium price logic.** The additional price sequence $\{\rho_t\}$ captures the scarcity value of storage capacity. When the capacity constraint binds, storage is scarce and $\rho_t > 0$: agents would like to carry more goods forward, but cannot. When the constraint does not bind, the scarcity rent disappears, and $\rho_t = 0$.

**What happens to interest rates and consumption dynamics?** Storage changes both (i) the feasibility set and (ii) the asset menu, so it can alter the equilibrium interest rate and the consumption pattern.

- **If storage is tight (small $\bar{x}$) or very costly (large $\delta$):** inventories cannot do much intertemporal smoothing. The economy remains close to the baseline exchange environment, so the interest rate and consumption can still display pronounced parity-driven oscillations.
- **If storage is abundant (large $\bar{x}$) and effective (small $\delta$):** inventories can absorb most of the endowment swings. Aggregate resources become smoother over time, so equilibrium consumption paths become smoother as well. Consequently, the need for large inter-agent bond trades is reduced.
- **Unlimited storage ($\bar{x} \to \infty$):** storage capacity is not scarce, so the competitive storage rent vanishes, $\rho_t \to 0$. Moreover, if storage can fully neutralize the aggregate timing differences in endowments, then the allocation can be implemented with minimal (or zero) net bond positions, $b_t \to 0$. In this limiting case, storage effectively replaces financial trades as the main vehicle for intertemporal smoothing.

## 2.5 Extension III: Uncertainty

Finally, suppose endowments are stochastic:

$$y_t^i = \bar{y}^i + \varepsilon_t^i, \qquad \varepsilon_t^i \sim \mathcal{N}(0,\sigma^2).$$

The household maximizes expected utility:

$$\max\ \mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \Big[u(c_{t+j}) + \lambda_{t+j}\big(y_{t+j} + b_{t+j-1}(1+r_{t+j-1}) - c_{t+j} - b_{t+j}\big)\Big].$$

The first-order conditions are

$$u'(c_t) = \lambda_t,$$

$$-\lambda_t + \mathbb{E}_t[\beta\lambda_{t+1}(1 + r_t)] = 0.$$

Combining these yields the stochastic Euler equation:

$$u'(c_t) = \beta(1 + r_t)\,\mathbb{E}_t\left[u'(c_{t+1})\right].$$

**Intuition.** Under uncertainty, optimal consumption equates marginal utility today to the expected discounted marginal utility tomorrow, adjusted by the real interest rate.

# 3 Stochastic Exchange and Dynamic Programming

## 3.1 Stochastic Exchange Environment

We now consider a stochastic pure exchange setting. Individual endowments are given by

$$y_t^i = \bar{y}^i + \varepsilon_t^i, \qquad \varepsilon_t^i \sim F(\varepsilon),$$

with shocks independently distributed over time.

The household Euler equation derived earlier generalizes to

$$u'(c_t) = \beta(1 + r_t)\,\mathbb{E}_t\left[u'(c_{t+1})\right].$$

This condition characterizes optimal intertemporal smoothing under uncertainty.

## 3.2 Dynamic Programming: A Review

Consider a generic infinite-horizon dynamic optimization problem:

$$\max_{\{c_t\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_t, s_t),$$

where:

- $c_t$ are control (choice) variables,
- $s_t$ are state variables.

State variables can be:

- *Exogenous* (e.g. productivity shocks),
- *Endogenous but predetermined* (e.g. skills or assets).

The state evolves according to the transition function

$$s_{t+1} = g(c_t, s_t).$$

### 3.3 The Bellman Equation

Define the value function as

$$V(s_0) = \max_{\{c_t\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_t, s_t) \quad \text{s.t.} \quad s_{t+1} = g(c_t, s_t).$$

The problem admits a recursive representation:

$$V(s) = \max_c \{u(c, s) + \beta V(s')\} \quad \text{with } s' = g(c, s).$$

Equivalently,

$$V(s) = \max_c \left\{u(c, s) + \beta V\big(g(c, s)\big)\right\}. \tag{Bellman}$$

### 3.4 Standard Results

1. There exists a unique value function $V$ solving the Bellman equation.
2. If $u(\cdot, s)$ is concave, then $V$ is concave.
3. Under regularity conditions, $V$ is differentiable.
4. $V$ can be approximated arbitrarily well by value function iteration.

### 3.5 Policy Functions

If $V$ is concave and $u$ is concave, the maximizer is unique. Hence the optimal policy is a function:

$$c = h(s).$$

### 3.6 First-Order Condition

Assuming interior solutions, the first-order condition is

$$\frac{\partial u(c, s)}{\partial c} + \beta \frac{\partial V(s')}{\partial s'} \frac{\partial g(c, s)}{\partial c} = 0, \quad s' = g(c, s).$$

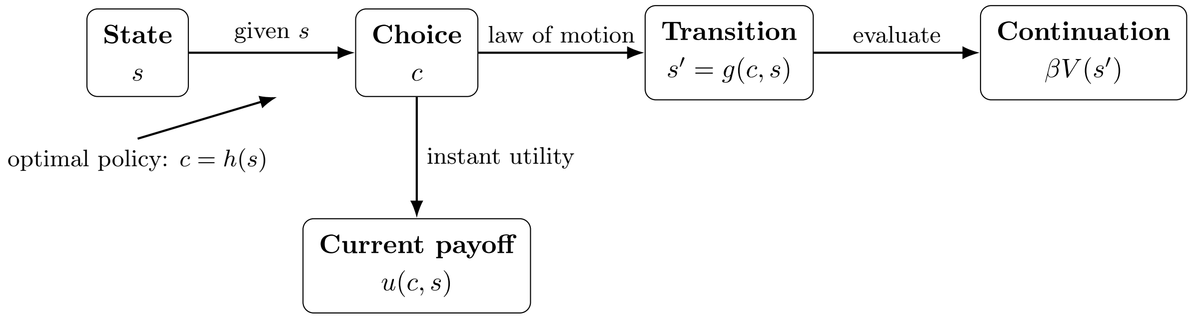

Figure 2: Recursive structure of the dynamic problem: current state $s$, control choice $c$, transition $s' = g(c, s)$, and continuation value $\beta V(s')$.

### 3.7 Envelope Condition

Using the Bellman equation evaluated at the optimal policy $c = h(s)$,

$$V(s) = u(h(s), s) + \beta V\big(g(h(s), s)\big),$$

differentiate both sides with respect to $s$. Applying the chain rule gives

$$\begin{aligned} V_s(s) &= u_c(h(s), s)\, h'(s) + u_s(h(s), s) \\ &\quad + \beta V_{s'}(s')\Big(g_c(h(s), s)\, h'(s) + g_s(h(s), s)\Big), \qquad s' = g(h(s), s). \end{aligned}$$

Collecting the terms that multiply $h'(s)$,

$$V_s(s) = u_s(h(s), s) + \beta V_{s'}(s')\, g_s(h(s), s) + \underbrace{\Big[u_c(h(s), s) + \beta V_{s'}(s')\, g_c(h(s), s)\Big]}_{=0}\, h'(s).$$

The bracketed term equals zero by the first-order condition for the choice of $c$ (evaluated at $c = h(s)$), so the derivative of the value function satisfies the envelope condition

$$V_s(s) = u_s(h(s), s) + \beta V_{s'}(s')\, g_s(h(s), s).$$

### 3.8 A Useful Trick: Choosing Next Period's State

A convenient technique is to rewrite the problem using next-period assets as the control.

**Example.** Let income follow

$$y_t = \bar{y} + \varepsilon_t, \qquad \varepsilon_t \sim F(\varepsilon),$$

and let preferences be $u(c)$. The budget constraint is

$$c_t + b_t = y_t + b_{t-1}(1 + r).$$

Define next-period assets:

$$A' \equiv b_t(1 + r).$$

Then the budget constraint becomes

$$c + \frac{A'}{1 + r} = y + A,$$

where the state variables are $(A, y)$.

### 3.9 Recursive Formulation with Assets

The Bellman equation can be written as

$$V(A, y) = \max_{c} \left\{u(c) + \beta\, \mathbb{E}\left[V(A', y')\right]\right\}.$$

Using the substitution $c = y + A - A'/(1 + r)$, we obtain the equivalent formulation:

$$V(A, y) = \max_{A'} \left\{u\left(y + A - \frac{A'}{1 + r}\right) + \beta\, \mathbb{E}\left[V(A', y')\right]\right\}.$$

### 3.10 Euler Equation via the Asset Choice

The first-order condition with respect to $A'$ is

$$-\frac{1}{1+r}u'(c) + \beta\,\mathbb{E}\left[\frac{\partial V(A', y')}{\partial A'}\right] = 0.$$

Applying the envelope condition,

$$\frac{\partial V(A, y)}{\partial A} = u'(c),$$

we obtain the Euler equation:

$$u'(c_t) = \beta(1+r)\,\mathbb{E}_t\left[u'(c_{t+1})\right].$$

**Intuition.** Writing the problem in terms of next-period assets simplifies the analysis: the Euler equation emerges directly from the FOC and the envelope condition, without explicitly differentiating the policy function.

# 4 Precautionary Savings and Stochastic Exchange Equilibrium

## 4.1 Concavity of the Value Function and Precautionary Motives

Assume that the value function $V(A, y)$ is strictly concave and differentiable in assets $A$. Then,

$$\frac{\partial^2 V}{\partial A^2} < 0.$$

In this case, the expected marginal value of future assets satisfies

$$\mathbb{E}\left[\frac{\partial V}{\partial A'}(A', \varepsilon')\right] > 0,$$

and is decreasing in $A'$.

The Euler equation can be written as

$$u'\left(\bar{y} + \varepsilon + A - \frac{A'}{1+r}\right) = \beta(1+r)\;\mathbb{E}\left[\frac{\partial V}{\partial A'}(A', \varepsilon')\right].$$

The left-hand side is increasing in $A'$ (since higher $A'$ reduces current consumption), while the right-hand side is decreasing in $A'$ due to concavity of $V$.

To see why an increase in $\mathrm{Var}(\varepsilon')$ can shift the RHS upward, fix $A'$ and define

$$\Phi(\varepsilon') \;\equiv\; V_{A'}(A', \varepsilon').$$

Consider a mean-preserving spread of $\varepsilon'$ that raises its variance while holding $\mathbb{E}[\varepsilon']$ constant. A second-order Taylor expansion of $\Phi$ around $\bar{\varepsilon} = \mathbb{E}[\varepsilon']$ yields

$$\mathbb{E}[\Phi(\varepsilon')] \;\approx\; \Phi(\bar{\varepsilon}) + \Phi'(\bar{\varepsilon})\;\mathbb{E}[\varepsilon' - \bar{\varepsilon}] + \frac{1}{2}\Phi''(\bar{\varepsilon})\;\mathbb{E}\left[(\varepsilon' - \bar{\varepsilon})^2\right].$$

Since $\mathbb{E}[\varepsilon' - \bar{\varepsilon}] = 0$, the linear term vanishes and

$$\mathbb{E}[\Phi(\varepsilon')] \;\approx\; \Phi(\bar{\varepsilon}) + \frac{1}{2}\Phi''(\bar{\varepsilon})\;\mathrm{Var}(\varepsilon').$$

Hence, if $\Phi$ is locally convex in $\varepsilon'$ (i.e. $\Phi''(\bar{\varepsilon}) > 0$), a higher variance increases the expected marginal

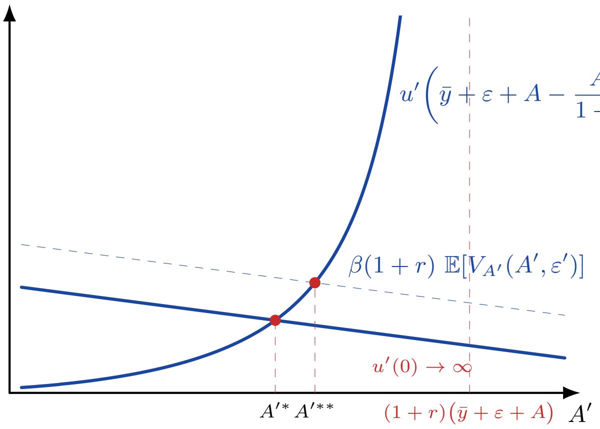

Figure 3: Precautionary-savings comparative static: a higher $\mathrm{Var}(\varepsilon')$ shifts the RHS upward and increases optimal saving $A'$.

value of future assets, shifting $\beta(1+r)\,\mathbb{E}[V_{A'}(A',\varepsilon')]$ upward. This convexity condition is closely related to *prudence*: in many standard savings problems, $V_{A'}$ is proportional to marginal utility of future consumption, and $u'''(c) > 0$ implies that marginal utility is convex, so mean-preserving spreads raise its expectation (a Jensen-type effect).

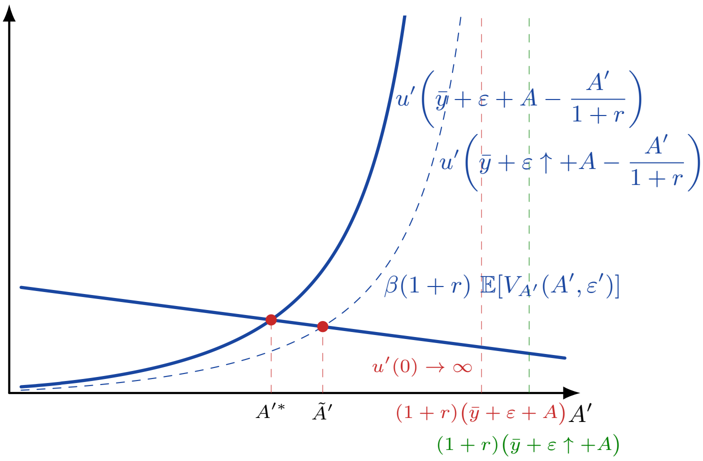

Figure 4: Current-income comparative static: a higher realized $\varepsilon$ shifts the LHS down and relaxes the borrowing limit, increasing optimal saving $A'$.

In the limit case where $u'(0)\to\infty$, the borrowing constraint never binds.

**Higher current income $\varepsilon$.** Fix $(A, r)$ and consider a higher realized shock $\varepsilon$ at date $t$. Holding $A'$ fixed, current consumption is

$$c \;=\; \bar{y}+\varepsilon+A-\frac{A'}{1+r},$$

so an increase in $\varepsilon$ raises $c$ one-for-one. Since $u'(\cdot)$ is decreasing, this lowers marginal utility for every $A'$, shifting the left-hand side

$$u'\left(\bar{y}+\varepsilon+A-\frac{A'}{1+r}\right)$$

*downward* as a function of $A'$ (the dashed curve in Figure 4).

At the same time, the higher current endowment relaxes the implicit borrowing limit. The constraint $c \geq 0$ implies

$$A' \;\leq\; (1+r)\big(\bar{y}+\varepsilon+A\big),$$

so a larger $\varepsilon$ moves the vertical asymptote to the right. Intuitively, higher current income both reduces the desire to save (lower marginal utility today) and expands feasible choices (more room to borrow).

Finally, the right-hand side $\beta(1+r)\,\mathbb{E}[V_{A'}(A',\varepsilon')]$ does not shift under this comparative static if the distribution of *future* shocks $\varepsilon'$ is unchanged (i.e. if we take expectations with respect to the same $F$). The new optimal choice $\tilde{A}'$ is therefore determined by the intersection of the unchanged RHS with the shifted LHS. In the figure, $\tilde{A}' > A'^{*}$: to restore the Euler condition after the downward shift in $marginalutility$, the agent chooses a higher $A'$, reducing current consumption and raising $u'(c)$ back to the level required by the continuation value.

## 4.2 Policy Functions

The optimal decision rules can be written as

$$c = h(A,\varepsilon), \qquad A' = s(A,\varepsilon),$$

where $\varepsilon \sim F(\varepsilon)$ with density $f(\varepsilon)$.

The Euler equation becomes

$$u'(c) = \beta(1+r)\int \frac{\partial u}{\partial c}\big(h(A',\varepsilon')\big)\, f(\varepsilon')\, d\varepsilon'.$$

Alternatively, using the budget constraint explicitly,

$$u'(c) = \beta(1+r)\int u'\left(\bar{y} + \varepsilon' + A' - \frac{s(A',\varepsilon')}{1+r}\right) f(\varepsilon')\, d\varepsilon'.$$

## 4.3 Stochastic Exchange Economy

Consider a continuum of individuals of unit mass. Individual endowments satisfy:

$$y_t(i) \sim f(y), \qquad \text{i.i.d. across individuals.}$$

Aggregate endowment is constant:

$$Y_t = \int_0^1 y_t(i)\, di = \bar{y}.$$

We conjecture a stationary equilibrium with a constant real interest rate:

$$1 + r_t = 1 + r.$$

## 4.4 Stationary Stochastic Equilibrium

**Definition.** A stationary stochastic equilibrium consists of:

- a constant real interest rate $1+r$,
- a policy function $s(A,y)$,
- an invariant distribution $g(A)$,

such that:

1. Given $1+r$, the policy function $s(A,y)$ solves the household problem.

2. Markets clear:
$$\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} s(A,y)\, f(y)\, g(A)\, dA\, dy = 0.$$
3. The invariant distribution $g(A)$ is implied by the policy function:
$$g(a) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} g(A)\, f(y)\, \mathbb{I}(a = s(A,y))\; dy\, dA, \quad \forall a,$$
where $\mathbb{I}(\cdot)$ is the indicator function.

This condition captures the endogenous evolution of the wealth distribution.

**Gráfico clave 2 (muy importante):** Plano $(A, y)$. Dibujar la política $s(A, y) = a$ como una curva decreciente en $(A, y)$. Marcar dos puntos $(A_1, y_1)$ y $(A_1, y_2)$ que mapean al mismo $a$. Mostrar con flechas cómo diferentes ingresos inducen la misma elección de activos futuros. Este gráfico ilustra cómo la política induce la distribución $g(A)$.

**Intuition.** Precautionary savings arise because uncertainty increases the marginal value of future assets. In equilibrium, the interest rate adjusts so that aggregate desired savings equal zero, while the distribution of wealth remains stationary over time.

# 5 Consumption Smoothing: Review

## 5.1 Permanent Income Hypothesis

Under the Permanent Income Hypothesis (PIH), consumption depends on permanent income:
$$c_t = f(y_t^P),$$
where permanent income $y_t^P$ is defined as the annuity value of expected lifetime resources. Observed income can be decomposed as
$$y_t = y_t^P + y_t^T,$$
where $y_t^T$ denotes transitory income.

## 5.2 Intertemporal Budget Constraint (Recall)

In the simple environment with a constant interest rate $r$,
$$c_0 + \frac{c_1}{1+r} + \frac{c_2}{(1+r)^2} + \cdots = y_0 + \frac{y_1}{1+r} + \frac{y_2}{(1+r)^2} + \cdots .$$

Thus, consumption choices reflect the present discounted value of income rather than its period-by-period realization.

## 5.3 The Random Walk Hypothesis (Hall, 1978)

Optimal consumption satisfies the Euler equation:
$$u'(c_t) = \beta(1 + r_t)\, \mathbb{E}_t \left[u'(c_{t+1})\right].$$

Under rational expectations, define the forecast error:

$$\mathbb{E}_t\left[u'(c_{t+1})\right] = u'(c_{t+1}) + \varepsilon_{t+1},$$

where

$$\mathbb{E}_t[\varepsilon_{t+1}] = 0, \qquad \mathrm{Cov}(\varepsilon_{t+1}, x_t) = 0 \quad \forall x_t \in \mathcal{I}_t,$$

with $\mathcal{I}_t$ denoting the information set at time $t$.

## 5.4 Implications

**1. Marginal Utility Follows a Random Walk.** Combining the Euler equation with the forecast error representation yields

$$u'(c_t) = \beta(1+r_t)\big(u'(c_{t+1}) + \varepsilon_{t+1}\big),$$

or equivalently,

$$u'(c_{t+1}) = \frac{1}{\beta(1+r_t)}u'(c_t) - \varepsilon_{t+1}.$$

Let $\gamma \equiv \frac{1}{\beta(1+r_t)}$. Then,

$$u'(c_{t+1}) = \gamma u'(c_t) + \varepsilon_{t+1}.$$

**Intuition.** Innovations to consumption reflect new information about lifetime resources. Bad news lowers the level of consumption permanently rather than inducing mean reversion.

**Gráfico clave:** trayectoria temporal de $u'(c_t)$ mostrando un random walk. Ante un shock negativo, la senda se desplaza hacia abajo de manera permanente, sin volver a la trayectoria original.

**2. Testable Euler Equation.** The Euler equation implies

$$u'(c_{t+1}) = \gamma u'(c_t) + \varepsilon_{t+1},$$

which provides a regression equation to estimate $\gamma$. Under rational expectations, OLS delivers an unbiased estimator.

## 5.5 Linearization

Using a first-order Taylor expansion,

$$u'(c_{t+1}) \approx u'(c_t) + u''(c_t)(c_{t+1} - c_t).$$

Substituting into the Euler equation implies that consumption growth follows:

$$c_{t+1} = \gamma c_t + \Gamma_1 x_t + \Gamma_2 y_t + \cdots + \eta_{t+1},$$

where $x_t$ includes all variables known at time $t$.

**Intuition.** Under the PIH and rational expectations, no variable in the information set should have predictive power for future consumption growth.

**Gráfico sugerido:** consumo $c_t$ en el tiempo con una senda suavizada; los shocks afectan el nivel, no la tasa de crecimiento esperada.

### 5.6 Empirical Implications

Hall (1978) finds that few observable variables significantly predict consumption growth, consistent with the Random Walk Hypothesis.

### 5.7 Key Assumptions

The empirical validity of the PIH and the Random Walk Hypothesis relies on:

1. Optimal consumption choice,
2. Rational expectations,
3. Aggregate (representative-agent) data,
4. Perfect credit markets,
5. Time-separable preferences.

Violations of these assumptions generate predictable movements in consumption.

### 5.8 Evidence and Extensions

**Wage Rigidity and Contracts.**

- Wilcox (1989): Social Security cost-of-living adjustments (COLAs).
- Shea (1995): Union wage contracts.

**Fiscal Policy.**

- Shapiro, Parker, and Slemrod (2006): 2001 U.S. tax rebates.

**Gráfico sugerido:** salario nominal $w$ con ajustes discretos en el tiempo (escalones), ilustrando rigideces contractuales y respuestas no suavizadas.

## 6 Hall (1988), Hansen–Singleton, and GMM Tests

### 6.1 Hall (1988): Lognormal Consumption and Interest-Rate Tests

Start from the standard Euler equation,

$$u'(c_t) = \beta(1 + r_t)\,\mathbb{E}_t\left[u'(c_{t+1})\right].$$

Assume CRRA preferences,

$$u(c) = \frac{c^{1-\frac{1}{\sigma}}}{1 - \frac{1}{\sigma}}, \qquad u'(c) = c^{-\frac{1}{\sigma}},$$

where $\sigma$ is the intertemporal elasticity of substitution (IES).

Then the Euler equation becomes

$$c_t^{-\frac{1}{\sigma}} = \beta(1 + r_t)\,\mathbb{E}_t\left[c_{t+1}^{-\frac{1}{\sigma}}\right].$$

**Distributional Assumption.** Hall (1988) assumes that, conditional on information at time $t$,

$$\ln c_{t+1} \sim \mathcal{N}(\mu_t, \sigma_c^2).$$

**Remark.** Although consumption is an endogenous variable, its *conditional* distribution is determined by the underlying structural shocks and the decision rule. In empirical work we typically do not observe all primitive shocks, so we impose a parsimonious parametric form on the conditional distribution of $\ln c_{t+1}$ (e.g. normality). This is a modelling shortcut: it is not a deep assumption about preferences, but a convenient way to evaluate $\mathbb{E}_t[c_{t+1}^{-1/\sigma}]$ in closed form and to obtain a regression equation we can estimate.

**Lognormality.** If a random variable $X$ is lognormal, then

$$\ln X \sim \mathcal{N}(\mu_X, \sigma_X^2) \quad \Rightarrow \quad \ln \mathbb{E}[X] = \mathbb{E}[\ln X] + \frac{1}{2} \operatorname{Var}(\ln X) = \mu_X + \frac{\sigma_X^2}{2}.$$

Apply this to

$$X = c_{t+1}^{-1/\sigma} = \exp\left(-\frac{1}{\sigma} \ln c_{t+1}\right).$$

Since $\ln c_{t+1} \sim \mathcal{N}(\mu_t, \sigma_c^2)$, we have

$$\ln X = -\frac{1}{\sigma} \ln c_{t+1} \sim \mathcal{N}\left(-\frac{\mu_t}{\sigma}, \frac{\sigma_c^2}{\sigma^2}\right).$$

Hence

$$\ln \mathbb{E}_t\left[c_{t+1}^{-1/\sigma}\right] = -\frac{\mu_t}{\sigma} + \frac{1}{2}\frac{\sigma_c^2}{\sigma^2}.$$

**Taking Logs of the Euler Equation.** Taking logs of the Euler condition,

$$-\frac{1}{\sigma} \ln c_t = \ln \beta + \ln(1 + r_t) + \ln \mathbb{E}_t\left[c_{t+1}^{-1/\sigma}\right],$$

and substituting the expression above yields

$$-\frac{1}{\sigma} \ln c_t = \ln \beta + \ln(1 + r_t) - \frac{\mu_t}{\sigma} + \frac{1}{2}\frac{\sigma_c^2}{\sigma^2}.$$

Multiplying by $-\sigma$ and recalling that $\mu_t = \mathbb{E}_t[\ln c_{t+1}]$,

$$\ln c_t = -\sigma \ln \beta - \sigma \ln(1 + r_t) + \mathbb{E}_t[\ln c_{t+1}] - \frac{1}{2}\frac{\sigma_c^2}{\sigma}.$$

Therefore,

$$\mathbb{E}_t[\ln c_{t+1}] - \ln c_t = \sigma \ln \beta + \sigma \ln(1 + r_t) + \frac{1}{2}\frac{\sigma_c^2}{\sigma}.$$

Using the approximation $\ln(1 + r_t) \approx r_t$ gives the testable prediction

$$\mathbb{E}_t[\Delta \ln c_{t+1}] = B + \theta r_t,$$

where

$$B \equiv \sigma \ln \beta + \frac{1}{2}\frac{\sigma_c^2}{\sigma}, \qquad \theta \equiv \sigma.$$

Under rational expectations,

$$\Delta \ln c_{t+1} = B + \theta r_t + \varepsilon_{t+1},$$

with $\varepsilon_{t+1}$ an innovation orthogonal to information at time $t$. This regression can be used to estimate the IES $\sigma$ from aggregate data.

## 6.2 Hansen and Singleton (1982): Multiple Assets and GMM

Now consider an environment with $N$ risky assets and a one-period risk-free bond. Let $a_{i,t}$ denote holdings of risky asset $i$, with price $p_{i,t}$ and dividend $d_{i,t}$. The household's budget constraint is

$$c_t + \sum_{i=1}^{N} p_{i,t} a_{i,t+1} + b_t = y_t + b_{t-1}(1+r_t) + \sum_{i=1}^{N} a_{i,t}(p_{i,t} + d_{i,t}).$$

The first-order condition for consumption is the same Euler equation as before,

$$u'(c_t) = \beta(1+r_t)\,\mathbb{E}_t[u'(c_{t+1})].$$

For each risky asset $i$, the FOC is

$$p_{i,t} u'(c_t) = \beta\ \mathbb{E}_t[(p_{i,t+1} + d_{i,t+1})u'(c_{t+1})]\,.$$

Define the gross return

$$R_{i,t+1} \equiv \frac{p_{i,t+1} + d_{i,t+1}}{p_{i,t}}.$$

Then the Euler equation can be written compactly as

$$1 = \beta\ \mathbb{E}_t\left[R_{i,t+1}\frac{u'(c_{t+1})}{u'(c_t)}\right], \qquad \forall i.$$

## 6.3 GMM Moment Conditions

Define the asset-specific moment condition

$$\varepsilon_{i,t+1}(\theta) \equiv \beta R_{i,t+1}\frac{u'(c_{t+1};\theta)}{u'(c_t;\theta)} - 1,$$

where $\theta$ collects preference parameters.

The model implies

$$\mathbb{E}\left[\varepsilon_{i,t+1}(\theta)\right] = 0 \qquad \forall i,$$

which can be estimated using GMM.

Specializing again to CRRA preferences,

$$u(c) = \frac{c^{1-\frac{1}{\sigma}}}{1-\frac{1}{\sigma}}, \qquad u'(c) = c^{-\frac{1}{\sigma}},$$

we obtain

$$\varepsilon_{i,t+1}(\theta) = \beta R_{i,t+1}\left(\frac{c_{t+1}}{c_t}\right)^{-\frac{1}{\sigma}} - 1,$$

with $\theta = (\sigma, \beta)$.

For the risk-free asset,

$$\varepsilon_{t+1}(\theta) = \beta(1+r_t)\left(\frac{c_{t+1}}{c_t}\right)^{-\frac{1}{\sigma}} - 1.$$

## 6.4 Gourinchas and Parker (2002): Life-Cycle Interpretation

Gourinchas and Parker (2002) apply similar tools to a life-cycle problem with non-stationary income and borrowing constraints. They estimate preference parameters by matching the predicted consumption profile over the life cycle to the data.

**Gráfico sugerido:** en el eje horizontal la edad $N$; en el eje vertical ingreso $y$ y consumo $c$. Dibujar una curva de ingreso con forma de joroba (crece, se estabiliza y luego cae) y una curva de consumo más suave y menos acentuada. A la derecha, anotar: "arranging consumption profiles over groups; no reliable individual longitudinal data".

## 6.5 Time-Inconsistent Preferences and "Irrational" Intertemporal Choice

The standard exponential discounting model assumes

$$U_t = \sum_{k=0}^{\infty} \beta^k u(c_{t+k}) = u(c_t) + \beta u(c_{t+1}) + \beta^2 u(c_{t+2}) + \cdots ,$$

where $\beta \in (0,1)$.

Under exponential discounting, preferences are *time-consistent*: the relative valuation between $t+k$ and $t+k+1$ does not depend on when the comparison is made.

### Evidence of Present Bias

Empirical work suggests that agents often display *present bias*. Important contributions include:

- DellaVigna and Malmendier (2006): evidence from gym membership contracts (per-day fee vs. annual membership) in the Boston Fitness Club.
- Laibson (1997, QJE): hyperbolic discounting.

### Quasi-Hyperbolic Discounting

A tractable alternative is the $\beta$–$\delta$ (quasi-hyperbolic) model:

$$U_t = u(c_t) + \delta \sum_{k=1}^{\infty} \beta^k u(c_{t+k}),$$

with $\delta \in (0,1]$ capturing present bias and $\beta \in (0,1)$ the standard long-run discount factor.

Explicitly,

$$U_t = u(c_t) + \delta \big[ \beta u(c_{t+1}) + \beta^2 u(c_{t+2}) + \cdots \big].$$

When $\delta < 1$, the agent places extra weight on current consumption relative to all future periods.

### Time Inconsistency

Consider the marginal rate of substitution (MRS) between $c_{t+k}$ and $c_{t+k+1}$.

**Comparison made at time $t$.**

$$\text{MRS}_t(t+k, t+k+1) = \frac{\delta \beta^k u'(c_{t+k})}{\delta \beta^{k+1} u'(c_{t+k+1})} = \frac{u'(c_{t+k})}{\beta u'(c_{t+k+1})}.$$

**Comparison made at time $t+k$.** Now $c_{t+k}$ is the present. The utility becomes

$$U_{t+k} = u(c_{t+k}) + \delta\big[\beta u(c_{t+k+1}) + \cdots\big],$$

so the MRS is

$$\text{MRS}_{t+k}(t+k, t+k+1) = \frac{u'(c_{t+k})}{\delta\beta u'(c_{t+k+1})}.$$

The extra $\delta$ term appears in the denominator.

Therefore,

$$\text{MRS}_t(t+k, t+k+1) \neq \text{MRS}_{t+k}(t+k, t+k+1) \quad \text{if } \delta < 1.$$

This is the source of *time inconsistency.*

Empirical estimates (e.g. Ainslie (1992), Ainslie and Haslam (1992)) suggest values around

$$\delta \approx \frac{2}{3}.$$

### Valuing Past Utility

An additional question concerns how individuals evaluate past utility.

Copin and Leahy (2004, JPE) study how agents retrospectively evaluate their life-cycle path. Empirically, individuals appear to evaluate past consumption with a discount structure that differs from the forward-looking one.

Conceptually, this generates a tension:

- At age 0, individuals may value the entire future path with $\beta = 1$.
- Later in life, evaluations of past consumption may follow a declining weighting scheme.

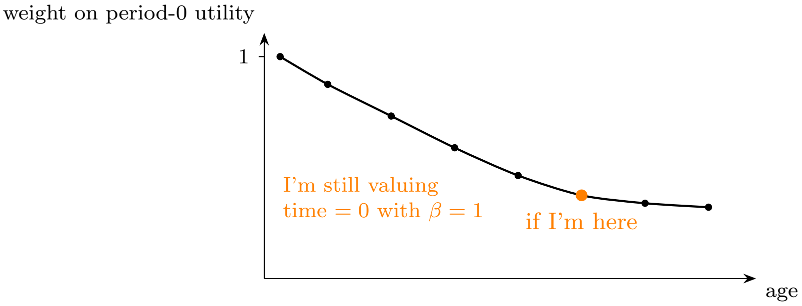

Figure 5: Retrospective weight on period-0 utility as a function of age.

The horizontal axis represents age, while the vertical axis measures the weight that the individual assigns *today* to the utility obtained in period 0. The downward-sloping black curve shows that, as the agent ages, the remembered utility from the distant past becomes progressively less important (the weight declines).

The orange point marks the individual's current age: even when the agent is far removed from period 0, a strictly positive weight is still assigned to that initial experience. The annotation "I'm still valuing time $=$ 0 with $\beta = 1$" illustrates a limiting case in which the agent, retrospectively, does not discount the past: the initial episode carries nearly as much weight as when it originally occurred.

The underlying idea is that temporal consistency requires the shape of this retrospective weighting curve to be coherent with the discounting rule used ex ante. If the ex post valuation of past utility does not align with the original discount structure, the agent would not be internally consistent when evaluating their own history of well-being.

The implied life-cycle profile is hump-shaped for income and smoother for consumption, but welfare evaluation requires consistency along the entire realized path.

In words: *one must live through the entire curve and be time-consistent in evaluating it.*

# 7 Labor Supply and Search

## 7.1 Neoclassical Labor Supply

Consider the standard static labor supply problem. The household chooses consumption $c$ and hours worked $N$ to maximize

$$\max_{c,N} \; u(c) - v(N)$$

subject to the budget constraint

$$pc = wN + A,$$

where $w$ is the nominal wage, $p$ the price level, and $A$ non-labor income.

**Functional Forms.** Assume CRRA utility over consumption and a convex disutility of labor:

$$u(c) = \frac{c^{1-\frac{1}{\sigma}}}{1-\frac{1}{\sigma}}, \qquad v(N) = \phi \frac{N^{1+\frac{1}{\eta}}}{1+\frac{1}{\eta}},$$

where:

- $\sigma$ is the intertemporal elasticity of substitution,
- $\eta$ is the Frisch labor supply elasticity,
- $\phi > 0$ scales the disutility of labor.

**First-Order Condition.** The optimality condition equates marginal disutility of labor to the real wage times marginal utility of consumption:

$$v'(N) = \frac{w}{p} u'(c).$$

Using the functional forms:

$$\phi N^{\frac{1}{\eta}} = \frac{w}{p} c^{-\frac{1}{\sigma}}.$$

**Log-Linearization.** Taking logs:

$$\log \phi + \frac{1}{\eta} \log N = \log\left(\frac{w}{p}\right) - \frac{1}{\sigma} \log c.$$

Rearranging:

$$\log N = \eta \log\left(\frac{w}{p}\right) - \frac{\eta}{\sigma} \log c + \text{const.}$$

Hence,

$$\frac{\partial \log N}{\partial \log(w/p)} = \eta.$$

**Definition** (Frisch Labor Supply Elasticity)**.** The Frisch elasticity $\eta$ measures the elasticity of hours worked with respect to the real wage, holding the marginal utility of wealth constant.

## 7.2 Empirical Specification

This leads to the regression equation

$$\ln N_t = \theta_0 + \theta_1 \log\left(\frac{w_t}{p_t}\right) + \theta_2 \log c_t + \varepsilon_t,$$

where:

$$\theta_1 = \eta, \qquad \theta_2 = -\frac{\eta}{\sigma}.$$

**Empirical Findings.**

- For single men and single women: $\eta \approx 0.1 - 0.2$ (very low elasticity).
- For married women, especially with children: $\eta \approx 0.3 - 1.0$ (more responsive).
- Typical estimates of $\theta_2$ range from $-0.05$ to $-0.25$.

**Remark.** It is important to distinguish between:

- *Intensive margin*: hours worked conditional on participation.
- *Extensive margin*: participation decision (entering or exiting the labor force).

Elasticities are generally larger along the extensive margin.

## 7.3 Labor Search Model

Consider a simple McCall-style search problem.

- An unemployed worker draws wage offers each period:

$$w \sim F(w), \quad \text{i.i.d. on } [0, B].$$

- Utility is linear in wages.
- Future utility is discounted at rate $\beta$.

**Decision.** Each period the worker decides:

- **Reject**: receive unemployment benefit $c$ and draw again next period.
- **Accept**: work forever at wage $w$, yielding present value

$$\frac{w}{1-\beta}.$$

**Interpretation.**

- $w$ includes both monetary wage and non-pecuniary job benefits (closeness to home, coworkers, childcare flexibility, etc.).
- $c$ can represent unemployment benefits or the value of leisure.

## 7.4 Labor Search and the Bellman Equation

Consider an unemployed worker who each period draws a wage offer

$$w \sim F(w), \qquad \text{i.i.d. on } [0, B].$$

Preferences are linear in wages and the worker discounts the future at rate $\beta$.

If the worker rejects the offer, she receives unemployment benefit $c$ and draws a new offer next period. If she accepts, she works forever at wage $w$, yielding lifetime utility

$$\frac{w}{1-\beta}.$$

Importantly, $w$ can be interpreted broadly: not only monetary compensation, but also non-monetary job characteristics (location, coworkers, amenities, etc.). Similarly, $c$ may represent unemployment insurance or the value of leisure.

## 7.5 Bellman Equation

The state variable is the current wage offer $w$. The Bellman equation is

$$V(w) = \max\left\{ \frac{w}{1-\beta},\ c + \beta\,\mathbb{E}[V(w')] \right\}.$$

The worker compares:

- Accepting the offer and receiving $\frac{w}{1-\beta}$,
- Rejecting and obtaining $c + \beta\,\mathbb{E}[V(w')]$.

## 7.6 Reservation Wage

There exists a reservation wage $\bar{w}$ such that:

$$V(w) = \begin{cases} \dfrac{w}{1-\beta}, & w \geq \bar{w}, \\ \dfrac{\bar{w}}{1-\beta}, & w \leq \bar{w}. \end{cases}$$

The reservation wage solves

$$v(\bar{w}) = \frac{\bar{w}}{1-\beta} = c + \beta\,\mathbb{E}[V(w')]. \tag{7.1}$$

Graphically, $V(w)$ is kinked and therefore not concave nor differentiable.

## 7.7 Characterizing the Reservation Wage

**Step 1: write $\mathbb{E}[V(w')]$ using the piecewise form.** Splitting the expectation at $\bar{w}$ gives

$$\mathbb{E}[V(w')] = \int_0^{\bar{w}} \frac{\bar{w}}{1-\beta}\, dF(w') + \int_{\bar{w}}^{B} \frac{w'}{1-\beta}\, dF(w').$$

Factor out $\frac{1}{1-\beta}$:

$$\mathbb{E}[V(w')] = \frac{1}{1-\beta} \left( \bar{w} \int_0^{\bar{w}} dF(w') + \int_{\bar{w}}^{B} w'\, dF(w') \right).$$

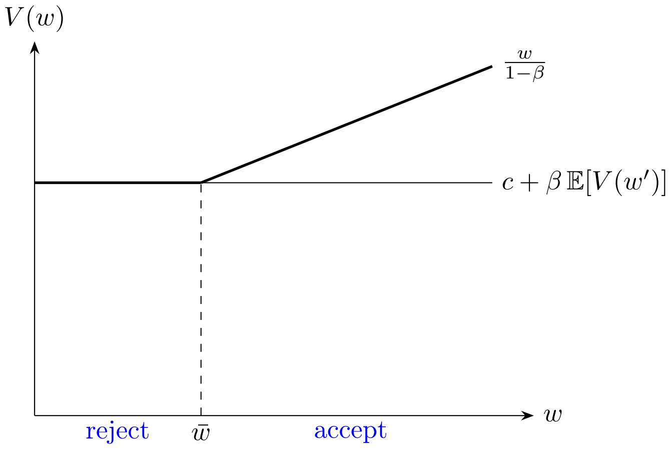

Figure 6: Value function with reservation wage.

Since $\int_0^{\bar{w}} dF(w') = F(\bar{w})$,

$$\mathbb{E}[V(w')] = \frac{1}{1-\beta}\left(\bar{w}F(\bar{w}) + \int_{\bar{w}}^{B} w'\,dF(w')\right).$$

**Step 2: substitute into the reservation-wage equation and clear denominators.** Plugging into $\frac{\bar{w}}{1-\beta} = c + \beta\,\mathbb{E}[V(w')]$ and multiplying both sides by $(1-\beta)$,

$$\bar{w} = c(1-\beta) + \beta\left(\bar{w}F(\bar{w}) + \int_{\bar{w}}^{B} w'\,dF(w')\right).$$

Rearranging,

$$\bar{w} - c(1-\beta) = \beta\bar{w}F(\bar{w}) + \beta\int_{\bar{w}}^{B} w'\,dF(w').$$

**Step 3: rewrite $\int_{\bar{w}}^{B} w'\,dF(w')$ using $\mathbb{E}[w']$ and integration by parts.** First, decompose the tail integral:

$$\int_{\bar{w}}^{B} w'\,dF(w') = \int_0^{B} w'\,dF(w') - \int_0^{\bar{w}} w'\,dF(w') = \mathbb{E}[w'] - \int_0^{\bar{w}} w'\,dF(w').$$

Now apply integration by parts (Stieltjes) to the last term. Take

$$u(w') = w', \qquad dv = dF(w'),$$

so that

$$du = dw', \qquad v = F(w').$$

Then

$$\int_0^{\bar{w}} w'\,dF(w') = \Big[w'F(w')\Big]_0^{\bar{w}} - \int_0^{\bar{w}} F(w')\,dw'.$$

Since $0 \cdot F(0) = 0$, this becomes

$$\int_0^{\bar{w}} w'\,dF(w') = \bar{w}F(\bar{w}) - \int_0^{\bar{w}} F(w')\,dw'.$$

Therefore,

$$\int_{\bar{w}}^{B} w' \, dF(w') = \mathbb{E}[w'] - \Big(\bar{w}F(\bar{w}) - \int_0^{\bar{w}} F(w') \, dw'\Big) = \mathbb{E}[w'] - \bar{w}F(\bar{w}) + \int_0^{\bar{w}} F(w') \, dw'.$$

**Step 4: substitute back and simplify.** Substituting into $\bar{w} - c(1-\beta) = \beta\bar{w}F(\bar{w}) + \beta \int_{\bar{w}}^{B} w' \, dF(w')$,

$$\begin{aligned}
\bar{w} - c(1-\beta) &= \beta\bar{w}F(\bar{w}) + \beta \left(\mathbb{E}[w'] - \bar{w}F(\bar{w}) + \int_0^{\bar{w}} F(w') \, dw'\right) \\
&= \beta\bar{w}F(\bar{w}) - \beta\bar{w}F(\bar{w}) + \beta\,\mathbb{E}[w'] + \beta \int_0^{\bar{w}} F(w') \, dw' \\
&= \beta\,\mathbb{E}[w'] + \beta \int_0^{\bar{w}} F(w') \, dw',
\end{aligned}$$

where the $\beta\bar{w}F(\bar{w})$ terms cancel.

$$\boxed{\bar{w} - c(1-\beta) = \beta\,\mathbb{E}[w'] + \beta \int_0^{\bar{w}} F(w') \, dw'.}$$

Define

$$g(\bar{w}) = \int_0^{\bar{w}} F(w') \, dw'.$$

Properties:

$$g(0) = 0, \qquad g(\bar{w}) \geq 0,$$

$$g'(\bar{w}) = F(\bar{w}) \in [0,1], \qquad g''(\bar{w}) = f(\bar{w}) \geq 0.$$

Thus $g$ is increasing and convex.

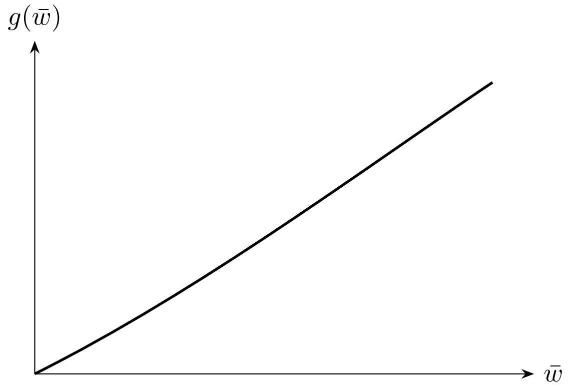

Figure 7: $g(\bar{w}) = \int_0^{\bar{w}} F(w')dw'$.

**Remark** (Interpretation of $g(0) = 0$)**.** Recall that

$$g(\bar{w}) = \int_0^{\bar{w}} F(w') \, dw'.$$

Evaluating at $\bar{w} = 0$ gives

$$g(0) = \int_0^{0} F(w') \, dw' = 0.$$

Economically, this means that if the reservation wage were zero, the worker would accept every wage offer. In that case, there is no region of rejection and therefore no option value from continued search.

The term $\beta g(\bar{w})$ in the reservation-wage equation

$$\bar{w} - c(1-\beta) = \beta\,\mathbb{E}[w] + \beta g(\bar{w})$$

captures precisely this option value of waiting. When $\bar{w} = 0$, that option disappears, so $g(0) = 0$.

## 7.8 Economic Intuition

The reservation wage balances:

- The immediate gain from accepting a job at wage $\bar{w}$,
- The option value of waiting for a better draw.

A higher unemployment benefit $c$ increases $\bar{w}$. A higher discount factor $\beta$ increases $\bar{w}$ because the worker becomes more patient and values future draws more. A mean-preserving increase in wage dispersion typically increases the option value of waiting and therefore raises $\bar{w}$.

The kink in $V(w)$ reflects a discrete choice: accept or reject. This generates non-differentiability and non-concavity, a key feature of search models.

## 7.9 Labor Search: Reservation Wage and Comparative Statics

**Bellman Equation.** State variable: current wage offer $w$.

$$V(w) = \max\left\{\frac{w}{1-\beta},\ c + \beta\mathbb{E}[V(w')]\right\}.$$

If the worker accepts the offer, she earns $w$ forever:

$$V^{\text{accept}}(w) = \frac{w}{1-\beta}.$$

If she rejects, she receives unemployment compensation $c$ and draws again next period:

$$V^{\text{reject}} = c + \beta\mathbb{E}[V(w')].$$

The reservation wage $\bar{w}$ solves

$$\frac{\bar{w}}{1-\beta} = c + \beta\mathbb{E}[V(w')].$$

Equivalently,

$$\bar{w} - c(1-\beta) = \beta\mathbb{E}[w'] + \beta\int_0^{\bar{w}} F(w')\,dw'.$$

The left-hand side (LHS) is linear in $\bar{w}$. The right-hand side (RHS) is increasing and concave (since $F$ is increasing).

**Properties:**

- The two curves always intersect (slope of RHS $< 1$ asymptotically).
- The intersection is unique.
- $V(w)$ is not concave and not differentiable at $\bar{w}$.

### Comparative Statics in $c$

If unemployment benefits increase ($c \uparrow$), the LHS shifts upward (parallel shift), increasing the reservation wage.

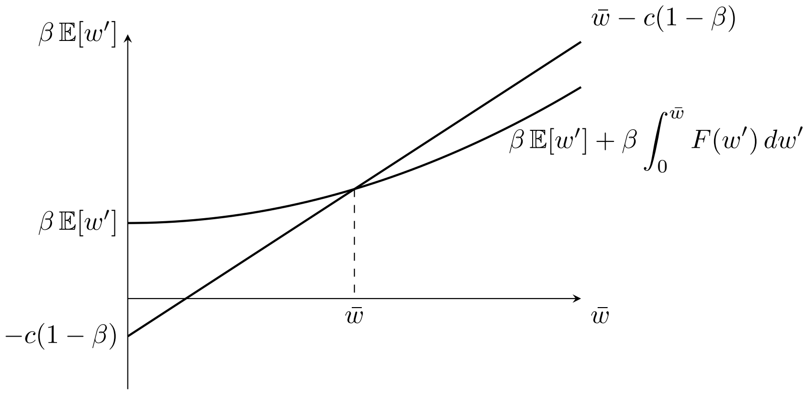

Figure 8: Reservation wage as intersection of LHS and RHS.

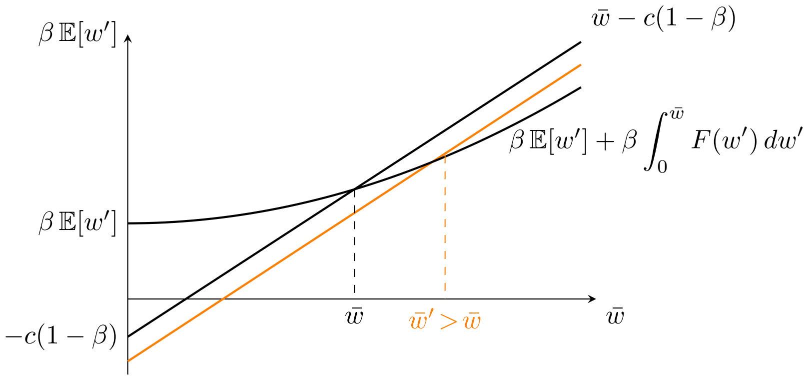

Figure 9: Increase in $c$ shifts the LHS down and raises the reservation wage $\bar{w}$ to $\bar{w}'$.

### Comparative Statics in $\beta$ (Patient Chooser)

More patient agents value future search opportunities more. An increase in $\beta$ shifts the RHS upward and also affects the intercept of the LHS.

It can be shown that[2]

$$\frac{\partial \bar{w}}{\partial \beta} > 0.$$

Moreover,

$$\lim_{\beta \to 1} \bar{w} = B,$$

where $B$ is the upper bound of the wage distribution.

---

[2] Proof. Define the implicit function

$$H(\bar{w}, \beta) = \bar{w} - c(1-\beta) - \beta\,\mathbb{E}[w] - \beta \int_0^{\bar{w}} F(w)\,dw.$$

The reservation wage satisfies $H(\bar{w}(\beta), \beta) = 0$. By the Implicit Function Theorem,

$$\frac{d\bar{w}}{d\beta} = -\frac{H_\beta}{H_{\bar{w}}}.$$

Direct differentiation yields

$$H_{\bar{w}} = 1 - \beta F(\bar{w}), \qquad H_\beta = c - \mathbb{E}[w] - \int_0^{\bar{w}} F(w)\,dw.$$

Using the reservation-wage equation to eliminate terms,

$$\frac{d\bar{w}}{d\beta} = \frac{\bar{w} - c}{\beta\big(1 - \beta F(\bar{w})\big)}.$$

Since $\beta > 0$ and $1 - \beta F(\bar{w}) > 0$, the derivative is strictly positive whenever $\bar{w} > c$.

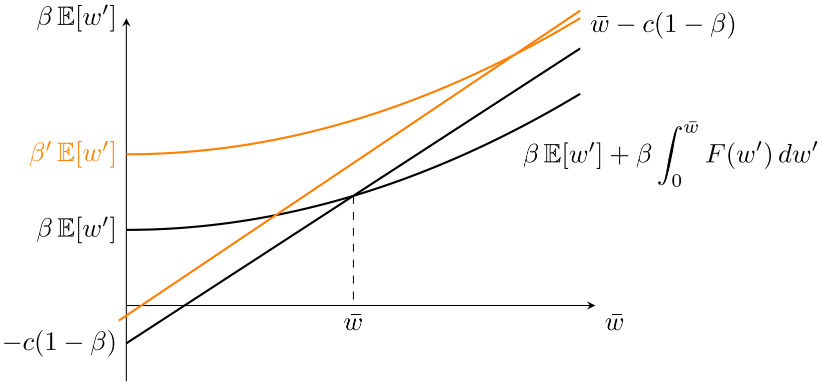

Figure 10: Effect of a higher $\beta$: both LHS and RHS shift up for a more patient chooser.

### Variance of Wage Offers and Mean-Preserving Spreads

Let $G$ be a mean-preserving spread (MPS) of $F$. A convenient characterization (second-order stochastic dominance) is

$$\int_0^x \big(G(w)-F(w)\big)\,dw \geq 0 \qquad \forall x \in [0,B],$$

with equality at $x = B$. Equivalently,

$$\int_0^x G(w)\,dw \;\geq\; \int_0^x F(w)\,dw \qquad \forall x.$$

Thus, for every cutoff $x$, the area under the CDF up to $x$ is *larger* under the more dispersed distribution. Importantly, an MPS keeps the mean fixed:

$$\mathbb{E}_G[w] = \mathbb{E}_F[w].$$

**Reservation wage equation.** The reservation wage satisfies

$$\bar{w} - c(1-\beta) = \beta\,\mathbb{E}[w] + \beta\int_0^{\bar{w}} F(w)\,dw.$$

Under an MPS, the mean $\mathbb{E}[w]$ is unchanged, so the effect operates through the term $\int_0^{\bar{w}} F(w)\,dw$. Since

$$\int_0^{\bar{w}} G(w)\,dw \;\geq\; \int_0^{\bar{w}} F(w)\,dw \quad \text{for every } \bar{w},$$

the RHS as a function of $\bar{w}$ shifts *upward* when the wage distribution becomes more dispersed.

Because the LHS,

$$\text{LHS}(\bar{w}) = \bar{w} - c(1-\beta),$$

is a straight increasing line, the unique intersection occurs at a strictly higher cutoff:

$$\bar{w}_G > \bar{w}_F.$$

**Economic intuition.** Rejecting an offer is an option. The worker receives

$$\mathbb{E}[\max\{w', \bar{w}\}]\,,$$

which is a convex function of $w'$. Under a mean-preserving spread, the expectation of any convex function increases (convex order / Jensen argument). Therefore the value of waiting rises.

Greater dispersion increases the upside potential (very high draws) while low draws can be rejected. The option value of search increases, so the worker becomes more selective and raises the reservation wage.

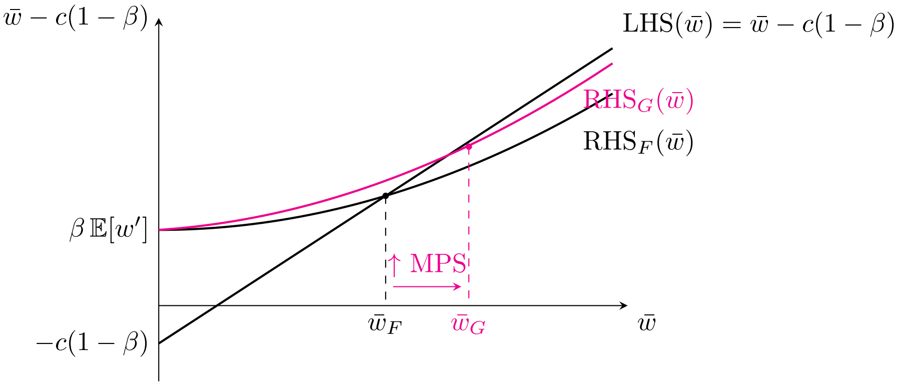

Figure 11: A mean-preserving spread shifts the RHS upward and raises the reservation wage.

## 7.10 Extensions of the Basic Search Model

### Allowing quits

Suppose that, after accepting a job, the worker is allowed to quit voluntarily and return to unemployment (and then resume search). In the basic stationary model, this option is never exercised in equilibrium: once a worker accepts a wage $w$, she either keeps the job forever or would have been better off rejecting in the first place.

Recall that the reservation wage $\bar{w}$ satisfies

$$v(\bar{w}) = \frac{\bar{w}}{1-\beta} = c + \beta \int v(w')\,dF(w'). \tag{6.3.7}$$

Consider a worker who currently holds an offer $w$. There are three mutually exclusive ways to respond:

**A1. Accept and keep the job forever.** Lifetime value:

$$A1(w) = \frac{w}{1-\beta}.$$

**A2. Accept and quit after $t$ periods.** The worker earns $w$ for $t$ periods and, starting in period $t$, returns to unemployment and behaves optimally thereafter. Thus,

$$A2(w,t) = \sum_{s=0}^{t-1} \beta^s w \;+\; \beta^t \Big(c + \beta \int v(w')\,dF(w')\Big).$$

Using (7.1), this becomes

$$A2(w,t) = \frac{w}{1-\beta} - \beta^t \frac{w-\bar{w}}{1-\beta}.$$

**A3. Reject the offer.** The worker takes unemployment compensation today and samples again next period:

$$A3 = \; c + \beta \int v(w')\,dF(w') = \frac{\bar{w}}{1-\beta}.$$

**Why quitting is dominated (intuition).** The quitting option $A2$ amounts to: *accept now, but later undo the acceptance and revert to the outside option.* The outside option available after quitting is exactly the continuation value of being unemployed, which (by definition of $\bar{w}$) is equivalent to holding an offer $\bar{w}$.

Because the environment is stationary, quitting after $t$ periods does not create a new opportunity that was unavailable at the time of the original decision; it only delays the moment at which the worker returns to the outside option. Hence:

- If $w < \bar{w}$, then accepting is already worse than rejecting. Moreover,

  $$A2(w,t) = A3 - \frac{(1-\beta^t)(\bar{w}-w)}{1-\beta} < A3, \qquad \text{and} \qquad A1(w) = \frac{w}{1-\beta} < A2(w,t).$$

  So $A1 \prec A2 \prec A3$.

- If $w > \bar{w}$, then accepting and keeping the job dominates rejecting. Also,

  $$A1(w) - A2(w,t) = \beta^t \frac{w-\bar{w}}{1-\beta} > 0, \qquad \text{so} \qquad A1 \succ A2.$$

  And since $A1(w) > \frac{\bar{w}}{1-\beta} = A3$, we have $A1 \succ A2 \sim A3$.

Therefore, for $w \neq \bar{w}$, the strategy "accept and quit later" is strictly dominated: when the offer is below the reservation wage one should reject immediately, and when it is above the reservation wage one should accept and keep the job. At $w = \bar{w}$, all three alternatives coincide.

### Firing Risk

Suppose that after the first period on the job, the worker faces probability $\alpha \in (0,1)$ of being fired each period. If fired, she spends one period unemployed receiving $c$ before drawing a new wage offer.

Let $\hat{v}(w)$ denote the value of holding offer $w$. Rejecting yields

$$\hat{v}^R = c + \beta E\hat{v}, \qquad E\hat{v} = \int \hat{v}(w')\, dF(w').$$

If the worker accepts, she receives $w$ today and:

- with probability $1-\alpha$, keeps the job next period (continuation value $\hat{v}(w)$);
- with probability $\alpha$, is fired, receives $c$ for one period, and then resumes search.

Thus,

$$\hat{v}^A(w) = w + \beta(1-\alpha)\hat{v}(w) + \beta\alpha\big[c + \beta E\hat{v}\big].$$

The Bellman equation is therefore

$$\hat{v}(w) = \max\left\{w + \beta(1-\alpha)\hat{v}(w) + \beta\alpha\big[c+\beta E\hat{v}\big],\ c + \beta E\hat{v}\right\}.$$

If the worker accepts,

$$\hat{v}(w) = w + \beta(1-\alpha)\hat{v}(w) + \beta\alpha\big[c+\beta E\hat{v}\big].$$

Rearranging,

$$\hat{v}(w)\big(1-\beta(1-\alpha)\big) = w + \beta\alpha\big[c+\beta E\hat{v}\big],$$

so[3]

$$\hat{v}^A(w) = \frac{w + \beta\alpha[c + \beta E\hat{v}]}{1 - \beta(1-\alpha)}.$$

Since $\hat{v}^A(w)$ is strictly increasing in $w$ while $\hat{v}^R$ is constant, the optimal policy is characterized by a reservation wage $\bar{w}$. Indifference at $w = \bar{w}$ implies[4]

$$\frac{\bar{w} + \beta\alpha[c + \beta E\hat{v}]}{1 - \beta(1-\alpha)} = c + \beta E\hat{v}.$$

Thus the reservation wage satisfies

$$\frac{\bar{w}}{1-\beta} = c + \beta \int \hat{v}(w')\, dF(w'). \tag{7.2}$$

- The reservation-wage equation has the same *form* as in the basic model, but with the new value function $\hat{v}(\cdot)$.
- Because jobs are now shorter-lived in expectation,

$$\hat{v}(w) < v(w) \quad \text{for all } w.$$

- Therefore, from (7.2) and the no-firing equation,

$$\bar{w}_{\text{firing}} < \bar{w}_{\text{no firing}}.$$

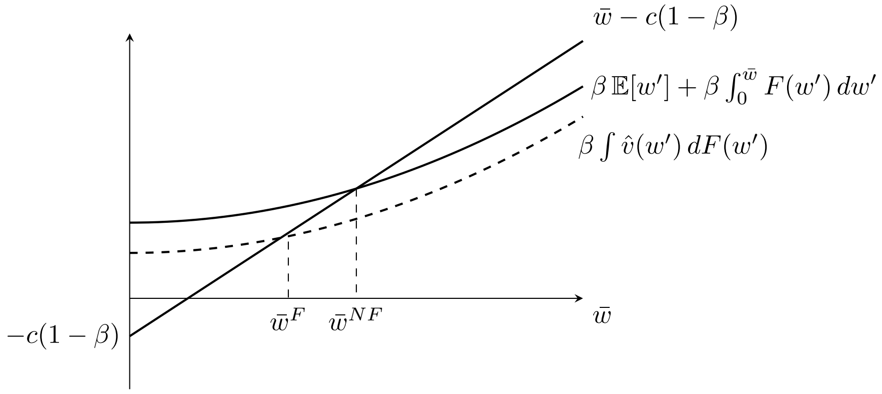

Figure 12: Reservation wage with and without firing risk. The RHS shifts downward under firings, reducing the reservation wage. I think that if we do the correct procedure, the intercept in the LHS should also change a bit. But the final intuition is the same.

When jobs may end exogenously, their expected duration falls. High wages are less valuable because they are less persistent. As a result, workers optimally lower their reservation wage: there is less reason to wait for very high-paying jobs when employment relationships are fragile.

---

[3] The denominator arises because employment continues only with probability $1 - \alpha$ each period. The value of a permanent job is therefore a perpetuity with effective discount factor $\beta(1-\alpha)$ rather than $\beta$.

[4] Multiplying both sides by $1 - \beta(1-\alpha)$ and canceling the common $\beta\alpha[c + \beta E\hat{v}]$ terms yields

$$\bar{w} = (1-\beta)\big(c + \beta E\hat{v}\big),$$

hence

$$\frac{\bar{w}}{1-\beta} = c + \beta E\hat{v}.$$

**Wage growth on the job and the reservation wage**

An unemployed worker draws an entry wage $w \sim F$ each period. If she accepts, the wage path is deterministic:

$$w_0 = w, \qquad w_t = \phi^t w \;\; (t \geq 1), \qquad \phi > 1,$$

and we assume $\beta\phi < 1$ so the present value is finite. If she rejects, she receives unemployment compensation $c$ for the period and draws a new offer next period. Accepting an offer with entry wage $w$ yields the discounted value

$$V_\phi^A(w) = \sum_{t=0}^{\infty} \beta^t (\phi^t w) = \frac{w}{1-\beta\phi}.$$

Relative to the basic model ($\phi = 1$), the only change is that the perpetuity factor becomes $\frac{1}{1-\beta\phi}$ instead of $\frac{1}{1-\beta}$. Let $V_\phi(w)$ be the value of holding offer $w$. Let $U_\phi$ be the continuation value of rejecting (the value of being unemployed after optimally behaving henceforth). Then

$$V_\phi(w) = \max\left\{\frac{w}{1-\beta\phi},\, U_\phi\right\}, \qquad U_\phi = c + \beta\,\mathbb{E}\left[V_\phi(w')\right].$$

As in the basic model, the optimal policy is a reservation rule. There exists $\bar{w}(\phi)$ such that the worker accepts iff $w \geq \bar{w}(\phi)$. The cutoff satisfies the indifference condition

$$\frac{\bar{w}(\phi)}{1-\beta\phi} = U_\phi.$$

If two economies differ only in the on-the-job wage growth rate, with $\phi_1 > \phi_2$ and $\beta\phi_i < 1$, then for every $w$,

$$V_{\phi_1}^A(w) = \frac{w}{1-\beta\phi_1} > \frac{w}{1-\beta\phi_2} = V_{\phi_2}^A(w).$$

Thus, holding fixed the outside option $U$, a higher $\phi$ makes *any* given entry wage more valuable, so the worker is willing to accept a lower entry wage. Formally, the acceptance payoff exhibits increasing differences in $(w, \phi)$, implying a monotone (downward) shift in the cutoff:

$$\boxed{\phi_1 > \phi_2 \;\Rightarrow\; \bar{w}(\phi_1) < \bar{w}(\phi_2).}$$

In the basic model, the entry wage $w$ is also the lifetime wage, so acceptance value is $w/(1-\beta)$. Here, a low entry wage can still be attractive because it is *leveraged* by future wage growth. A higher $\phi$ increases the slope of the acceptance value in $w$, so the worker becomes less selective about entry wages and the reservation wage falls.

### 7.10.1 Exploding Offers (Offers Stay Valid for $k$ Periods)

**Idea.** In the baseline model, the worker either accepts the *current* offer or rejects it and loses it forever. An "exploding offer" generalization assumes that an offer remains available for a limited number of periods. In the notes, the offer lasts for four periods, so the worker can remember the current offer and the last three offers, and can accept the best among them.

**State.** Let the state be the vector of available offers

$$\mathbf{w} \equiv (w_0, w_1, w_2, w_3),$$

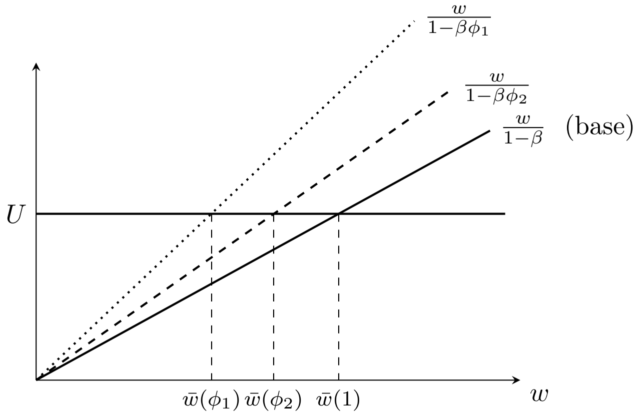

Figure 13: Wage growth raises the value of accepting any entry wage. A higher $\phi$ steepens the acceptance value $w/(1-\beta\phi)$, lowering the reservation wage.

where $w_0$ is the current-period offer and $w_1, w_2, w_3$ are the offers received in the previous three periods that have not yet expired.

**Bellman equation.** If the worker accepts, she chooses the best available offer and works forever at that wage, giving payoff $\max_{i\in\{0,1,2,3\}} w_i/(1-\beta)$. If she rejects, she gets unemployment payoff $c$ today and tomorrow the offer set "shifts" and a new offer is drawn. A convenient representation is

$$V(w_0, w_1, w_2, w_3) = \max\left\{\frac{\max\{w_0, w_1, w_2, w_3\}}{1-\beta},\; c + \beta\,\mathbb{E}\Big[V(w_0', w_0, w_1, w_2)\Big]\right\},$$

where $w_0' \sim F$ is next period's new draw (and $w_3$ drops out because it expires).

**Intuition.** This extension increases the *option value* of waiting: rejecting does not mean you lose all good opportunities, because a high offer can be "banked" for a few periods. As a result, the reservation wage typically increases relative to the baseline: the worker can afford to be more selective because the continuation value of rejecting is higher.

### 7.10.2 Job Destruction

**Idea.** Once employed, the job can end exogenously with probability $\delta$ each period ("job destruction"). This forces us to distinguish the value of being unemployed from the value of being employed at a given wage.

**Two value functions.** Let $V^U$ denote the value of unemployment and $V^E(w)$ the value of employment at wage $w$. If employed at wage $w$, the worker gets $w$ today; next period the job survives with probability $1-\delta$ and is destroyed with probability $\delta$ (returning the worker to unemployment):

$$V^E(w) = w + \beta\Big[(1-\delta)V^E(w) + \delta V^U\Big].$$

Solving this fixed point yields

$$V^E(w) = \frac{w + \beta\delta V^U}{1-\beta(1-\delta)}.$$

When unemployed, the worker receives $c$ and draws a new offer $w' \sim F$ next period. Upon observing

$w'$, she can accept (and become employed at $w'$) or reject (remain unemployed):

$$V^U = c + \beta\, \mathbb{E}\Big[\max\{V^E(w'), V^U\}\Big].$$

**Reservation wage.** Define the reservation wage $\bar{w}$ by indifference:

$$V^E(\bar{w}) = V^U.$$

Using the closed form for $V^E(w)$, this becomes a single equation in $\bar{w}$ (implicitly via the expectation in $V^U$).

**Remark.** In the notes, there is a warning "this is not right" next to an expression that tries to write the employed value as something like $w + \beta\{\delta\, \mathbb{E}[V^U(w')] + (1-\delta)V^E(w)\}$ while mixing states and expectations. The clean fix is exactly to separate the states ($U$ vs. $E$), write two Bellman equations, and only take expectations over *next period's* randomness (e.g. next period's offer when unemployed, and next period's destruction event when employed).

**Intuition.** Job destruction lowers the value of accepting, because employment is no longer an absorbing state. Equivalently, it raises the "effective discounting" of future wage flows. All else equal, this tends to lower the reservation wage relative to the no-destruction case (because the prize from waiting for a very high permanent job is smaller when jobs can end).

### 7.10.3 New Careers vs. New Jobs

**Motivation.** The notes distinguish two layers of heterogeneity: a *career* component and a *job* component. The logic is:

> A new career requires a new job, but a new job does not necessarily require a new career.

**Payoff decomposition.** Let the flow payoff ("wage") be

$$w = \theta + \varepsilon,$$

where $\theta$ is the career payoff (persistent across jobs within that career) and $\varepsilon$ is the job-specific payoff (i.i.d. across jobs). In the notes, $\theta$ has distribution $F$ and $\varepsilon$ has distribution $G$.

**Bellman equation with two search margins.** Let the state be $(\theta, \varepsilon)$: your current career draw and your current job draw within that career. A natural three-way choice is:

1. *Stay* in the current job: receive $\theta + \varepsilon$ forever (given the baseline "work forever" structure).
2. *Search for a new job within the same career:* keep $\theta$, draw a new $\varepsilon' \sim G$.
3. *Switch career:* draw a new $\theta' \sim F$ and (since a new career implies a new job) also draw a new $\varepsilon' \sim G$.

A compact way to write this is

$$V(\theta,\varepsilon) = \max\left\{ \underbrace{\theta + \varepsilon + \beta V(\theta,\varepsilon)}_{\text{stay}},\ \underbrace{\theta + \int \Big(\varepsilon' + \beta V(\theta,\varepsilon')\Big)\, dG(\varepsilon')}_{\text{new job, same career}},\ \underbrace{\iint \Big(\theta' + \varepsilon' + \beta V(\theta',\varepsilon')\Big)\, dG(\varepsilon')\, dF(\theta')}_{\text{new career (and job)}} \right\}.$$

The first term can be rewritten in closed form as

$$\theta+\varepsilon+\beta V(\theta,\varepsilon) \;\Rightarrow\; V(\theta,\varepsilon)=\frac{\theta+\varepsilon}{1-\beta} \quad \text{if it is optimal to stay forever.}$$

**Collapsing the two "search" objects (as in the notes).** The notes then name the two expectation terms:

$$\Gamma(\theta) \;\equiv\; \theta+\int \Big(\varepsilon'+\beta V(\theta,\varepsilon')\Big)\, dG(\varepsilon') \qquad \text{("new job, same career").}$$

This is *not* a constant in general; it depends on $\theta$.

The "new career (and job)" value is an expected value, i.e. a number:

$$q \;\equiv\; \iint \Big(\theta'+\varepsilon'+\beta V(\theta',\varepsilon')\Big)\, dG(\varepsilon')\, dF(\theta') \qquad \text{("new career and new job").}$$

With this notation, the Bellman equation becomes

$$V(\theta,\varepsilon)=\max\left\{\frac{\theta+\varepsilon}{1-\beta},\ \Gamma(\theta),\ q\right\}.$$

**Cutoff for staying.** From

$$\frac{\theta+\varepsilon}{1-\beta}\geq \max\{\Gamma(\theta),q\} \quad\Longleftrightarrow\quad \theta+\varepsilon\geq(1-\beta)\max\{\Gamma(\theta),q\},$$

the notes define the (possibly $\theta$-dependent) cutoff

$$\hat{\varepsilon}(\theta) \;\equiv\; (1-\beta)\max\{\Gamma(\theta),q\}-\theta,$$

so that

$$\varepsilon\geq\hat{\varepsilon}(\theta) \quad\Longrightarrow\quad \text{stay.}$$

**Career-retention threshold and the piecewise form for $\hat{\varepsilon}(\theta)$.** The notes then compare $\Gamma(\theta)$ to $q$ and consider the locus

$$\Gamma(\theta)=q.$$

There is a threshold $\hat{\theta}$ ("large enough to keep your career") such that for $\theta\geq\hat{\theta}$ it is optimal to *keep* the career (i.e. the relevant outside option is $\Gamma(\theta)$ rather than $q$).

A key claim made in the notes is:

$$\text{If } \theta>\hat{\theta}, \text{ then } \hat{\varepsilon}(\theta) \text{ is a constant (independent of } \theta).$$

In contrast, in the region where switching careers is the relevant outside option,

$$\Gamma(\theta)\leq q \quad\Longrightarrow\quad \hat{\varepsilon}(\theta)=(1-\beta)q-\theta.$$

**Intuition and "why this is interesting".** This extension cleanly separates two option values: (i) the option value of sampling *jobs* (high $\varepsilon$ draws) while holding career fit fixed, and (ii) the option value of sampling *careers* (high $\theta$ draws), which is typically a rarer, larger, more persistent improvement. It also clarifies why policies or frictions that affect mobility can have very different effects depending on whether they operate on the job margin (e.g. low-cost job switching within an industry) or the career margin (e.g. costly retraining).

**Curiosity.** With $(\theta, \varepsilon)$, the continuation value can become even more non-concave and non-differentiable than in the scalar-wage model, because the max operator now compares objects that differ in which component is being resampled. Graphically, one often gets multiple regions in $(\theta, \varepsilon)$ space: "stay", "switch job", or "switch career".

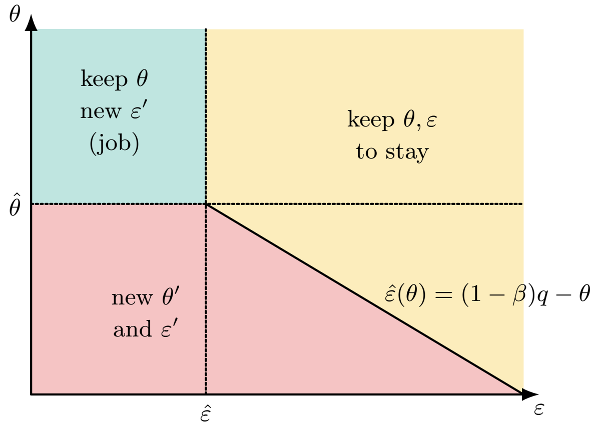

Figure 14: State-space partition into "switch job", "stay", and "switch career" regions, as in the lecture sketch.

**Interpretation of the lower triangular region.** The lower-right triangular region corresponds to states with relatively low $\theta$ (a "bad" career) but sufficiently high $\varepsilon$ (an exceptionally good job match). In this region,

$$\frac{\theta + \varepsilon}{1-\beta} \geq q,$$

so the value of staying exceeds the value of switching to a new career (which yields the expected value $q$). The intuition is that even if the persistent career component $\theta$ is weak, a sufficiently high job-specific draw $\varepsilon$ raises the present value of remaining in the current position. Switching careers would require giving up this unusually good job match and redrawing both $\theta'$ and $\varepsilon'$ from their distributions. Thus, the agent rationally stays: it is optimal to keep a very good job even within a mediocre career.

# 8 Investment Spending (Demand)

## 8.1 National Income Identity and Why Investment Matters

We start from the expenditure identity

$$GDP_t = C_t + I_t + G_t + NX_t.$$

Consumption is the largest component of GDP and tends to be relatively smooth over time. The investment component, by contrast, is much more volatile and therefore central for understanding business-cycle fluctuations.

**Remark.** In U.S. data, investment is typically around 15–20% of GDP, but its short-run volatility is disproportionately large relative to its average share.

## 8.2 Main Investment Categories in the NIPA Accounts

A useful decomposition is:

- **Residential investment** (housing structures),
- **Non-residential investment** (structures, equipment/software, intellectual property),
- **Change in private inventories**.

Even though inventories are a small share of GDP in levels, they can explain a relevant fraction of high-frequency GDP movements because inventory investment is very volatile.

**Intuition.** Small components can have large contributions to fluctuations if they move a lot from one period to the next.

## 8.3 Definitions: Capital, Investment, and Depreciation

**Definition. Physical capital** is the stock of productive assets (plants, factories, machinery, equipment).

**Definition. Investment** is spending on newly produced capital goods that adds to the productive stock.

**Definition. Depreciation** is the wear and tear (or obsolescence) that reduces the effective capital stock over time.

## 8.4 Stock–Flow Distinction and Capital Accumulation

Capital is a *stock* variable; investment is a *flow* variable. In standard discrete time, capital evolves according to

$$k_{t+1} = (1-\delta)k_t + I_t,$$

where $\delta \in (0,1)$ is the depreciation rate.

- $(1-\delta)k_t$: undepreciated capital carried into period $t+1$,
- $I_t$: new investment that augments next period's productive capacity.

## 8.5 Measurement Notes and Practical Classification Issues

In practice, boundaries across subcategories (e.g., structures vs. equipment) are not always sharp. National-account classifications are ultimately statistical conventions implemented by agencies such as the BEA. For theory, what matters is to keep track of the economic object (stock vs. flow), while for empirics one must respect the accounting definitions used in the data.

## 8.6 Business-Cycle Relevance of Housing

Residential investment is often one of the earliest and most cyclical components of aggregate demand. Because housing is interest-sensitive and forward-looking, it frequently acts as a leading indicator in business-cycle dynamics.

## 8.7 Illustrative Composition (Class Example: 2018)

As an empirical benchmark discussed in class, total investment was about 18% of GDP, with an approximate decomposition:

- residential investment: 3.8% of GDP,
- non-residential investment: 12.4% of GDP,

- change in inventories: 0.1% of GDP.

A further non-residential split used in lecture was:

- structures: 2.7% of GDP,
- equipment and software: 5.7% of GDP,
- intellectual property products: 4.1% of GDP.

**Remark.** These shares are useful for calibration and intuition: investment is smaller than consumption in levels, but much more volatile over the business cycle.

### 8.8 Why Inventories Matter for Fluctuations

Although inventory investment is small on average, it can account for a substantial fraction of short-run GDP fluctuations because it is very volatile and strongly procyclical.

## 9 User Cost of Capital and Investment Demand

### 9.1 From Hiring Labor to Hiring Capital Services

For labor, the textbook static condition is

$$p_t^y \cdot MPN_t = w_t \qquad \Longleftrightarrow \qquad MPN_t = \frac{w_t}{p_t^y}.$$

Analogously, if the firm rents capital services in a spot market,

$$p_t^y \cdot MPK_t = R_t,$$

where $R_t$ is the real rental cost of one unit of capital services.

**Intuition.** In both cases the firm compares a marginal benefit (extra output valued at goods prices) with a marginal cost (the input price).

### 9.2 Owner-User Capital and the No-Arbitrage Logic

Suppose instead the firm owns capital. Let $p_t^k$ be the purchase price of one unit of capital good. At date $t$, one unit of resources worth $p_t^k$ can be used in two ways:

1. buy capital and operate it for one period, yielding current marginal product plus resale value,
2. save in financial markets at gross return $(1 + r_t)$.

No-arbitrage implies

$$(1 + r_t)p_t^k = p_t^y MPK_t + (1 - \delta)p_{t+1}^k.$$

Equivalently,

$$p_t^y MPK_t = (1 + r_t)p_t^k - (1 - \delta)p_{t+1}^k.$$

## 9.3 User Cost Formula with Capital Gains

Rewrite the previous expression as

$$p_t^y MPK_t = p_t^k \left[ r_t + \delta - (1-\delta)\frac{p_{t+1}^k - p_t^k}{p_t^k} \right].$$

Define capital-goods inflation

$$\pi_{t+1}^k \equiv \frac{p_{t+1}^k - p_t^k}{p_t^k}.$$

Then

$$p_t^y MPK_t = p_t^k \left[ r_t + \delta - (1-\delta)\pi_{t+1}^k \right].$$

**Remark.** The term $-(1-\delta)\pi_{t+1}^k$ is the capital-gains correction: when the capital good is expected to appreciate, the effective user cost falls because resale value is higher.

If capital-goods prices are expected to remain constant, $p_{t+1}^k = p_t^k$, then there are no expected capital gains from holding capital. In that case, the value of the marginal product of capital must equal the *user cost* of capital:

$$p_t^y MPK_t = (r_t + \delta)p_t^k.$$

The right-hand side has two components: the financial opportunity cost of using resources to buy capital rather than saving them, $r_t p_t^k$, and the loss in value due to depreciation, $\delta p_t^k$. Thus, a unit of capital is worth employing only if the additional value it produces covers both the forgone return and the depreciation cost.

## 9.4 Clear Comparative Statics

Holding other terms fixed:

- a higher real interest rate $r_t$ raises the user cost of capital,
- a higher depreciation rate $\delta$ raises the user cost,
- higher expected capital-goods inflation $\pi_{t+1}^k$ lowers the user cost.

Hence desired capital demand is downward sloping in user cost and, in particular, negative in $r_t$ under standard assumptions on technology.

## 9.5 Continuous-Time Primer (Same Economic Condition)

Consider

$$\max_{\{I(t)\}_{t\geq 0}} \int_0^\infty e^{-rt} \Big[ F(k(t)) - p^k(t)I(t) \Big] dt$$

subject to

$$\dot{k}(t) = I(t) - \delta k(t).$$

Using the current-value Hamiltonian

$$\mathcal{H} = F(k) - p^k I + \lambda\big(I - \delta k\big),$$

the first-order condition with respect to investment is

$$\frac{\partial \mathcal{H}}{\partial I} = -p^k(t) + \lambda(t) = 0,$$

so that

$$\lambda(t) = p^k(t).$$

The costate equation is

$$\dot{\lambda}(t) = r\lambda(t) - \frac{\partial \mathcal{H}}{\partial k} = r\lambda(t) - \big(F_k(k(t)) - \delta\lambda(t)\big).$$

Rearranging,

$$F_k(k(t)) = (r+\delta)\lambda(t) - \dot{\lambda}(t).$$

Substituting $\lambda(t) = p^k(t)$ yields

$$F_k(k(t)) = (r+\delta)p^k(t) - \dot{p}^k(t) = p^k(t)\left(r + \delta - \frac{\dot{p}^k(t)}{p^k(t)}\right).$$

**Intuition.** The firm equates the marginal product of installed capital to its user cost: the financing cost plus depreciation, net of expected capital gains on the capital good.

## 9.6 Discrete-Time Bellman Version

The analogous discrete-time problem is

$$\sum_{t=0}^{\infty} \left(\frac{1}{1+r_t}\right)^t \left[F(k_t) - p_t^k I_t\right]$$

subject to

$$k_{t+1} = (1-\delta)k_t + I_t.$$

Using the law of motion to substitute out investment,

$$I_t = k_{t+1} - (1-\delta)k_t,$$

the Bellman equation is

$$V_t(k_t) = \max_{k_{t+1}} \left\{ F(k_t) - p_t^k \left[k_{t+1} - (1-\delta)k_t\right] + \frac{1}{1+r_t} V_{t+1}(k_{t+1}) \right\}.$$

The first-order condition is

$$-p_t^k + \frac{1}{1+r_t} V'_{t+1}(k_{t+1}) = 0,$$

or equivalently,

$$V'_{t+1}(k_{t+1}) = (1+r_t)p_t^k.$$

The envelope condition is

$$V'_t(k_t) = F_k(k_t) + (1-\delta)p_t^k.$$

Evaluating the envelope condition one period ahead gives

$$V'_{t+1}(k_{t+1}) = F_k(k_{t+1}) + (1-\delta)p_{t+1}^k.$$

Combining this expression with the first-order condition yields

$$(1+r_t)p_t^k = F_k(k_{t+1}) + (1-\delta)p_{t+1}^k.$$

Hence,

$$F_k(k_{t+1}) = (1 + r_t)p_t^k - (1 - \delta)p_{t+1}^k,$$

which is exactly the discrete-time no-arbitrage condition underlying the user cost of capital.

## 9.7 From Desired Capital to Investment Demand

In the frictionless neoclassical model, once the firm has identified its desired capital stock, investment is pinned down by the law of motion for capital. Abstracting from capital gains and normalizing the price of capital goods to one, the static optimality condition is

$$F_k(k_{t+1}^*) = r_t + \delta,$$

where $k_{t+1}^*$ denotes the desired next-period capital stock.

Since capital evolves according to

$$k_{t+1} = (1 - \delta)k_t + I_t,$$

investment must satisfy

$$I_t = k_{t+1}^* - (1 - \delta)k_t.$$

Thus, once the desired capital stock is determined, this expression defines an investment demand function.

**Intuition.** The neoclassical investment problem can be separated into two steps. First, determine the desired capital stock from the condition that marginal product equals user cost. Second, choose investment so that the capital stock moves toward that desired level.

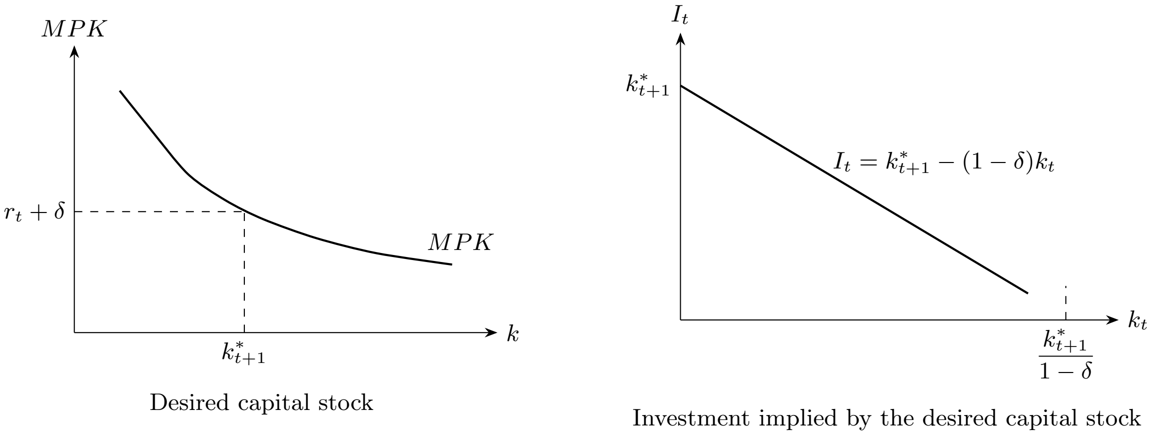

Desired capital stock

Investment implied by the desired capital stock

Figure 15: From desired capital to investment demand.

## 9.8 Tobin's $q$ Theory of Investment

A different route to investment demand is provided by Tobin's $q$-theory. Define average Tobin's $q$ as

$$q_t^{\text{avg}} \equiv \frac{\text{Market value of the firm}_t}{\text{Replacement cost of capital}_t}.$$

In empirical work, the numerator is often approximated by the market value of debt plus equity, while the denominator is measured using book value or an estimate of replacement cost. The basic prediction of the theory is simple: investment should be an increasing function of $q$. If the market values installed capital more than its replacement cost, the firm has an incentive to invest and expand its capital stock.

**Remark.** Average Tobin's $q$ is not, in general, the same object as the shadow value relevant for optimal investment. The variable that enters the firm's first-order condition is *marginal* $q$, not average $q$. Hayashi's contribution was to show that, under restrictive but economically important conditions, the two coincide.

## 9.9 Why Might $q \neq 1$?

In a frictionless Modigliani–Miller environment with no rents, no intangible capital, and no adjustment costs, one would expect the market value of firms to be closely tied to the replacement cost of their capital, so that $q$ would be near one.

In practice, however, $q$ can differ from one for several reasons:

- intangible capital is omitted from measured replacement cost,
- monopoly power generates rents that raise firm value,
- adjustment costs make installed capital more valuable than uninstalled capital.

This last mechanism is central in modern investment theory.

## 9.10 Adjustment Costs and Marginal $q$

Following Tobin's insight, a standard dynamic formulation introduces adjustment costs explicitly. The firm solves

$$\max_{\{I(t)\}_{t\geq 0}} \int_0^\infty e^{-rt}\Big[F(k(t)) - p^k(t)I(t) - c(I(t),k(t))\Big]dt$$

subject to

$$\dot{k}(t) = I(t) - \delta k(t).$$

The term $c(I(t),k(t))$ denotes the firm's *adjustment cost* of investment. This is distinct from the direct purchase cost of capital, $p^k(t)I(t)$. While $p^k(t)I(t)$ captures the market price of acquiring new capital goods, $c(I(t),k(t))$ captures the additional cost of transforming investment expenditures into installed and productive capital.

Economically, adjustment costs reflect the fact that expanding the capital stock is typically not frictionless. Installing new machines, reorganizing production, training workers, and temporarily disrupting existing operations all make rapid investment costly. Thus, $c(I(t),k(t))$ represents the extra resource cost associated with changing the firm's capital stock.

It is natural for this cost to depend both on investment, $I(t)$, and on the existing capital stock, $k(t)$. Higher investment usually implies higher adjustment costs, while the same amount of investment may be easier for a large firm to absorb than for a small one. For this reason, adjustment costs are often modeled as depending on the investment rate, $I(t)/k(t)$, rather than on investment alone.

The main role of adjustment costs in the model is to prevent the firm from adjusting its capital stock instantaneously. Without such costs, the firm would jump immediately to its desired capital level. With adjustment costs, investment becomes gradual, and the shadow value of installed capital—marginal $q$—plays a central role in determining the optimal pace of adjustment.

Using the current-value Hamiltonian[5],

$$\mathcal{H} = F(k) - p^k I - c(I,k) + q\big(I - \delta k\big),$$

the first-order condition with respect to investment is

$$\frac{\partial \mathcal{H}}{\partial I} = -p^k(t) - c_I(I(t), k(t)) + q(t) = 0,$$

so

$$q(t) = p^k(t) + c_I(I(t), k(t)).$$

If we normalize the price of capital goods to one, $p^k(t) = 1$, then

$$q(t) = 1 + c_I(I(t), k(t)).$$

The costate equation is

$$\dot{q}(t) = rq(t) - \frac{\partial \mathcal{H}}{\partial k},$$

that is,

$$\dot{q}(t) = rq(t) - \Big[F_k(k(t)) - c_k(I(t), k(t)) - \delta q(t)\Big].$$

Equivalently,

$$rq(t) - \dot{q}(t) = F_k(k(t)) - c_k(I(t), k(t)) - \delta q(t).$$

**Intuition.** Marginal $q$ is the shadow value of installed capital. Without adjustment costs, one extra unit of installed capital is worth exactly its purchase price. With adjustment costs, installed capital is more valuable because expanding the capital stock is itself costly.

## 9.11 Economic Meaning of the Adjustment-Cost Function

The function $c(I,k)$ captures internal costs of changing the capital stock. These are not purchase costs of capital goods themselves, but installation, organizational, and disruption costs associated with adjusting productive capacity.

A standard set of assumptions is:

$$c_I(I,k) > 0 \quad \text{for sufficiently high investment,}$$

$$c_I(I,k) < 0 \quad \text{for sufficiently low investment,}$$

and

$$c_{II}(I,k) \geq 0.$$

The last condition is convexity: large changes in investment are disproportionately costly.

---

[5] The variable $q(t)$ is the current-value co-state variable associated with the law of motion for capital, $\dot{k}(t) = I(t) - \delta k(t)$. Formally, it is the dynamic analogue of a Lagrange multiplier in a static constrained optimization problem. Economically, $q(t)$ measures the shadow value of an additional unit of installed capital at time $t$, that is, the marginal contribution of an extra unit of $k(t)$ to the value of the firm. For this reason, $q(t)$ is interpreted as marginal $q$.

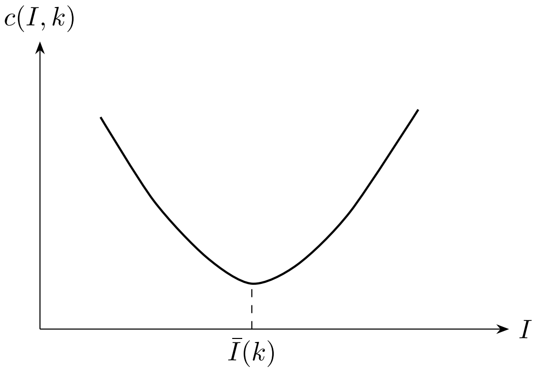

Figure 16: Adjustment costs are typically modeled as convex in investment around a benchmark level $\bar{I}(k)$.

## 9.12 Hayashi's Homogeneous Adjustment-Cost Specification

A particularly useful specification, emphasized by Hayashi, is

$$c(I,k) = k\,\phi\!\left(\frac{I}{k}\right),$$

where $\phi$ is convex.

This formulation is attractive because adjustment costs depend on the investment rate $I/k$, not on the absolute scale of the firm. A common example is

$$c(I,k) = k\left(\frac{I}{k} - \delta\right)^2.$$

Under this specification, replacement investment $I = \delta k$ is costless, while faster or slower adjustment is costly.

If $p^k(t) = 1$, the first-order condition becomes

$$q = 1 + c_I(I,k).$$

Since

$$c_I(I,k) = \frac{\partial}{\partial I}\left[k\,\phi\!\left(\frac{I}{k}\right)\right] = \phi'\!\left(\frac{I}{k}\right),$$

we obtain

$$q = 1 + \phi'\!\left(\frac{I}{k}\right).$$

**Remark.** This equation is the key investment relation in $q$-theory. Marginal $q$ exceeds one precisely when the marginal adjustment cost of investment is positive.

Because $\phi$ is convex, $\phi'$ is increasing. Therefore the map from $I/k$ to $q$ is monotone and can be inverted:

$$\phi'\!\left(\frac{I}{k}\right) = q - 1 \qquad \Longrightarrow \qquad \frac{I}{k} = h(q),$$

with $h'(q) > 0$.

Hence the investment rate is an increasing function of marginal $q$.

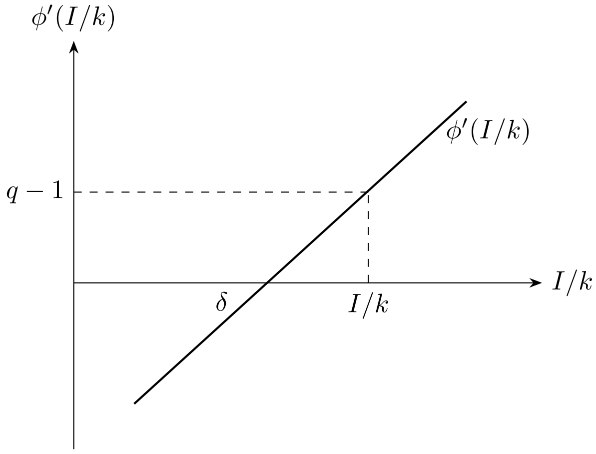

Figure 17: Convex adjustment costs imply a monotone relation between marginal $q$ and the investment rate.

## 9.13 Quadratic Example

Suppose

$$\phi(x) = (x-\delta)^2.$$

Then

$$\phi'(x) = 2(x-\delta),$$

so the investment equation becomes

$$q = 1 + 2\left(\frac{I}{k} - \delta\right).$$

Solving for the investment rate,

$$\frac{I}{k} = \delta + \frac{1}{2}(q-1).$$

This linear specification is especially convenient because the investment rate depends on a single state variable, $q$, and responds positively to it.

**Intuition.** If $q = 1$, the firm merely replaces depreciated capital, so $I/k = \delta$. If $q > 1$, installed capital is more valuable at the margin than its replacement cost, and the firm invests faster than replacement. If $q < 1$, investment falls below replacement.

## 9.14 Average Tobin's $Q$ versus Marginal $q$

It is crucial to distinguish between two related but conceptually different objects.

First, *marginal* $q$ is the shadow value of one additional unit of installed capital:

$$q_t \equiv \frac{\partial V_t}{\partial k_t}.$$

Second, *average Tobin's* $Q$ is the ratio of the market value of the firm to the replacement cost of its capital stock. When the price of capital goods is normalized to one, this can be written as

$$Q_t^{\text{Tobin}} \equiv \frac{V_t}{k_t}.$$

In general, these two objects are not the same:

$$q_t \neq Q_t^{\text{Tobin}}.$$

This distinction matters because the firm's optimal investment decision is governed by *marginal* $q$, not by average $Q$. The first-order condition for investment depends on the shadow value of an additional unit of capital, whereas financial market data more naturally provide information about the average valuation of the firm.

**Remark.** The theory of investment is written in terms of marginal $q$, but empirical work usually observes only a proxy for average Tobin's $Q$. This gap is one of the central challenges in bringing $q$-theory to the data.

Hayashi's central insight is that, under a restrictive but important set of conditions, these two objects coincide. In particular, if

- production exhibits constant returns to scale,
- firms are price takers in competitive markets,
- the adjustment-cost function is homogeneous of degree one in $(I, k)$,
- there is a single type of capital,
- and there are no financing constraints,

then

$$Q_t^{\text{Tobin}} = q_t.$$

This result is what makes empirical $Q$-investment regressions theoretically meaningful: observed market valuation can then be interpreted as a sufficient statistic for the firm's incentive to invest.

## 9.15 A Convenient Quadratic Adjustment-Cost Specification

A standard quadratic specification is

$$c(I, k) = \frac{\theta}{2} k \left( \frac{I}{k} - \delta \right)^2.$$

Equivalently,

$$c(I, k) = k\,\phi\!\left(\frac{I}{k}\right), \qquad \phi(x) = \frac{\theta}{2}(x - \delta)^2.$$

Then

$$c_I(I, k) = \theta \left( \frac{I}{k} - \delta \right),$$

so under the normalization $p^k = 1$, the investment first-order condition

$$q = 1 + c_I(I, k)$$

becomes

$$q = 1 + \theta \left( \frac{I}{k} - \delta \right).$$

Solving for the investment rate,

$$\frac{I}{k} = \delta + \frac{1}{\theta}(q - 1).$$

**Intuition.** The parameter $\theta$ governs how costly it is to change the capital stock rapidly. A larger $\theta$ means a steeper adjustment-cost schedule and therefore a weaker response of investment to a given change in $q$.

**Remark.** Some notes and papers instead write the quadratic cost as

$$c(I,k) = \frac{k}{2\theta}\left(\frac{I}{k} - \delta\right)^2.$$

That alternative normalization is equivalent, but then $\theta$ measures responsiveness rather than the curvature of adjustment costs. The economics is the same; only the parameterization changes.

## 9.16 Why a Simple Investment Regression Is Not So Innocent

The relation

$$\frac{I_{ft}}{K_{ft}} = \delta + \frac{1}{\theta}(q_{ft} - 1)$$

might suggest estimating an equation of the form

$$\frac{I_{ft}}{K_{ft}} = \beta_0 + \beta_1(q_{ft} - 1) + e_{ft}.$$

But this raises two immediate issues.

First, what is $q_{ft}$ empirically? The theory is written in terms of marginal $q$, while the data usually provide an imperfect proxy based on market valuation, namely average Tobin's $Q$.

Second, one cannot simply append a residual without further assumptions. Since $q_{ft}$ is a forward-looking endogenous shadow value, the orthogonality conditions needed for regression interpretation are nontrivial. Measurement error, omitted variables, and financing frictions all matter.

## 9.17 The Dynamic Equation for Marginal $q$

Recall the current-value Hamiltonian

$$\mathcal{H} = F(k) - I - c(I,k) + q(I - \delta k),$$

where we have normalized $p^k = 1$.

The costate equation is

$$\dot{q}(t) = rq(t) - \mathcal{H}_k.$$

Since

$$\mathcal{H}_k = F_k(k) - c_k(I,k) - \delta q,$$

it follows that

$$F_k(k) - c_k(I,k) - \delta q = rq - \dot{q}.$$

Now suppose

$$c(I,k) = k\,\phi\left(\frac{I}{k}\right).$$

Then, letting $x \equiv I/k$,

$$c_k(I,k) = \phi(x) - x\phi'(x).$$

Therefore,

$$F_k(k) - \phi\left(\frac{I}{k}\right) + \frac{I}{k}\phi'\left(\frac{I}{k}\right) = -\dot{q} + q(r+\delta).$$

Define

$$\xi\left(\frac{I}{k}\right) \equiv \phi\left(\frac{I}{k}\right) - \frac{I}{k}\phi'\left(\frac{I}{k}\right).$$

Then the law of motion for marginal $q$ can be written compactly as

$$F_k(k(t)) - \xi\left(\frac{I(t)}{k(t)}\right) = -\dot{q}(t) + (r+\delta)q(t),$$

or equivalently,

$$\dot{q}(t) - (r+\delta)q(t) = \xi\left(\frac{I(t)}{k(t)}\right) - F_k(k(t)).$$

**Remark.** This is a linear differential equation in $q(t)$. It makes explicit that marginal $q$ summarizes the discounted value of future marginal profitability of capital, net of the marginal adjustment-cost wedge.

Using the integrating factor $e^{-(r+\delta)t}$, and imposing the standard no-bubbles / transversality condition, one obtains

$$q(t) = \int_t^\infty e^{-(r+\delta)(s-t)} \left[F_k(k(s)) - \xi\left(\frac{I(s)}{k(s)}\right)\right] ds.$$

**Intuition.** Marginal $q$ is forward-looking. It is high when an additional unit of installed capital raises future profits by a lot, and low when the gains from capital accumulation are small or when adjustment frictions are severe.

## 9.18 Hayashi's Result: When Does Average $Q$ Equal Marginal $q$?

The appeal of $q$-theory is that, under strong conditions, one variable is enough to summarize investment incentives.

Hayashi's key result is that average Tobin's $Q$ equals marginal $q$ when the following conditions hold:

1. the production technology is constant returns to scale,
2. markets are competitive, so firms are price takers,
3. there are no financing constraints,
4. there is a single type of capital,
5. the adjustment-cost function is homogeneous of degree one in $(I, k)$.

Under these assumptions,

$$Q_t^{\text{Tobin}} = q_t.$$

**Remark.** This is what makes empirical $Q$-investment regressions theoretically attractive: observable market valuation can then stand in for the shadow value relevant for optimal investment.

Examples of constant-returns-to-scale production functions include

$$F(K, N) = AK,$$

$$F(K, N) = AK^\alpha N^{1-\alpha},$$

and more generally production functions in which the exponents sum to one. A constant-returns CES example is

$$F(K, N) = A\left(\omega K^{\rho} + (1-\omega)N^{\rho}\right)^{1/\rho}.$$

## 9.19 Taxes, Adjustment Costs, and Investment Demand

To connect the theory more directly with empirical work, it is useful to specialize the firm's problem to the case in which the price of capital goods is normalized to one and adjustment costs take the homogeneous form $c(I,k) = k\phi(I/k)$. With a proportional corporate tax, the firm solves

$$\max_{\{I(t)\}_{t\geq 0}} \int_0^{\infty} e^{-rt}\Big[(1-\tau(t))\Pi(k(t)) - I(t) - k(t)\phi\Big(\frac{I(t)}{k(t)}\Big)\Big]dt$$

subject to

$$\dot{k}(t) = I(t) - \delta k(t).$$

Here $\tau(t)$ is the corporate tax rate, $\Pi(k(t))$ denotes operating profits before investment and adjustment costs, and $k(t)\phi(I(t)/k(t))$ is the adjustment-cost term. This specification is convenient because it makes the cost of adjustment depend on the investment rate $I(t)/k(t)$, rather than on investment in levels.

Using the current-value Hamiltonian,

$$\mathcal{H} = (1-\tau(t))\Pi(k(t)) - I(t) - k(t)\phi\Big(\frac{I(t)}{k(t)}\Big) + q(t)\big(I(t) - \delta k(t)\big),$$

the first-order condition with respect to investment is

$$\frac{\partial \mathcal{H}}{\partial I} = -1 - \phi'\Big(\frac{I(t)}{k(t)}\Big) + q(t) = 0.$$

Hence,

$$q(t) = 1 + \phi'\Big(\frac{I(t)}{k(t)}\Big).$$

If $\phi'' > 0$, this condition can be inverted to obtain

$$\frac{I(t)}{k(t)} = h\big(q(t)\big),$$

for some increasing function $h(\cdot)$. Thus, as in the standard $q$-theory, the shadow value of installed capital is the sufficient statistic for the firm's investment rate.

**Intuition.** The firm invests up to the point at which the shadow value of an additional unit of installed capital, $q(t)$, equals the purchase cost of investment plus the marginal adjustment cost of installing it. Once adjustment costs are introduced, higher $q$ induces more investment, but the response is gradual rather than instantaneous.

The co-state equation implies

$$(r+\delta)q(t) - \dot{q}(t) = (1-\tau(t))\Pi_k(k(t)) - \phi\Big(\frac{I(t)}{k(t)}\Big) + \frac{I(t)}{k(t)}\phi'\Big(\frac{I(t)}{k(t)}\Big).$$

This expression shows how taxes enter the model: a higher corporate tax reduces the after-tax marginal profitability of capital and therefore tends to reduce $q$, which in turn lowers investment demand.

## 9.20 Average $Q$, Marginal $q$, and Empirical Implementation

The theoretical object governing investment is marginal $q$, that is, the shadow value of an additional unit of installed capital. Empirically, however, researchers typically observe average Tobin's $Q$, defined as the market value of the firm relative to the replacement cost of its capital stock.

Under the conditions of Hayashi's theorem—constant returns to scale, a single type of capital, perfect competition, and adjustment costs homogeneous of degree one in $(I, k)$—average $Q$ coincides with marginal $q$. In that case, the investment equation can be written in terms of observable valuation measures:

$$\frac{I_{ft}}{K_{ft}} = \beta_0 + \beta_1 Q_{ft} + u_{ft}.$$

A common interpretation comes from the quadratic adjustment-cost case, in which the investment rule is approximately linear in $q$. Then a small estimated $\beta_1$ corresponds to a steep marginal adjustment-cost schedule and therefore to sluggish investment responses.

**Remark.** This is the origin of the small-$q$ coefficient puzzle: empirically, the estimated response of investment to $Q$ is often much weaker than the simple model predicts.

The same logic extends to the taxed version of the model. Taxes change the law of motion for $q$ by altering after-tax marginal profitability, but they do not overturn the basic investment-demand relation $\frac{I}{K} = h(q)$. Under the same homogeneity conditions, one can still use observed $Q$ as an empirical proxy for the relevant shadow value.

## 9.21 Cash Flow Regressions and the Financing-Constraints Interpretation

A prominent empirical specification augments the investment equation with cash flow:

$$\frac{I_{ft}}{K_{ft}} = \beta_0 + \beta_1 Q_{ft} + \beta_2 \frac{\text{Cash Flow}_{ft}}{K_{ft}} + u_{ft}.$$

In the basic $q$-theory, $Q$ should be the sufficient statistic for investment once it correctly measures marginal $q$. Therefore, under the benchmark model,

$$\beta_2 = 0.$$

If instead $\beta_2 > 0$, one possible interpretation is that internal funds matter for investment because firms face financing frictions. This is the logic behind the classic cash-flow sensitivity literature associated with Fazzari, Hubbard, and Petersen.

Their findings suggest that cash flow remains strongly significant even after controlling for $Q$, while the coefficient on $Q$ itself often remains small. This seems difficult to reconcile with the simplest version of $q$-theory and motivates the view that financial constraints may affect firms' investment decisions.

## 9.22 Measurement Error, Identification, and the Critique of Cash-Flow Regressions

The financing-constraints interpretation is not, however, uncontroversial. A central concern is that measured $Q$ may be a noisy proxy for the true marginal $q$. If so, the empirical significance of cash flow need not imply that internal finance has an independent causal effect on investment.

A useful way to organize the issue is to introduce a latent profitability component $f_{it}$ driving the firm's true investment opportunities:

$$\frac{I_{it}}{K_{it}} = \alpha_0 + \alpha_1 f_{it} + \nu_{it},$$

while observed valuation and cash flow satisfy

$$Q_{it} = f_{it} + \varepsilon_{it}^{Q},$$

$$\frac{\text{Cash Flow}_{it}}{K_{it}} = \gamma_0 + \gamma_1 f_{it} + \gamma_x X_{it} + \varepsilon_{it}^{C}.$$

In this representation, both $Q_{it}$ and cash flow contain information about the same underlying fundamental $f_{it}$. If $Q_{it}$ is measured with error, then regressing investment on observed $Q_{it}$ generates attenuation bias in the estimate of $\beta_1$. Once cash flow is added, it may pick up part of the variation in $f_{it}$ that measured $Q$ fails to capture. As a result, the cash-flow coefficient may appear significant even in the absence of financing constraints.

**Intuition.** If the econometrician observes a noisy proxy for marginal $q$, then the regression may understate the role of valuation and attribute explanatory power to cash flow that actually belongs to omitted fundamentals. In this sense, cash flow may be acting as a proxy for true investment opportunities rather than as direct evidence of financing frictions.

This is the core identification problem in the cash-flow literature. It also explains why simply "controlling for cash flow" is conceptually delicate: cash flow may contain independent information about financing conditions, but it may also proxy for expected profitability that is imperfectly measured by $Q$.

## 9.23 Fundamental $q$ and Later Empirical Work

One response to the measurement-error critique is to construct a measure of *fundamental* $q$ from the discounted value of expected future marginal profits. The basic idea is to recover the shadow value of installed capital from firm fundamentals rather than from stock-market valuation alone.

Under a simplified version of the model in which adjustment costs depend only on investment, $c(I,k) = c(I)$, the investment rule still takes the form

$$I(t) = h\big(q(t)\big).$$

If, for expositional simplicity, one also abstracts from depreciation, $\delta = 0$, the co-state equation reduces to

$$\dot{q}(t) = rq(t) - (1-\tau(t))\Pi_k(k(t)).$$

Forward iteration then implies that $q(t)$ can be written as the discounted value of future after-tax marginal profitability:

$$q(t) = \int_t^{\infty} e^{-r(s-t)}(1-\tau(s))\Pi_k(k(s))\,ds.$$

In discrete time, the same logic yields a representation of the form

$$q_{it}^{\text{fund}} = \sum_{j=0}^{\infty} \left(\frac{1-\delta}{1+r}\right)^{j} E_t[(1-\tau_{i,t+j})\Pi_{k,i,t+j}]\,.$$

In a production representation, $\Pi_k$ may be replaced by $F_k$, the marginal product of capital.

This approach, associated with later work such as Gilchrist and Himmelberg, aims to build a cleaner empirical proxy for marginal $q$. If investment responds more strongly to fundamental $q$ than to stock-market $Q$, then at least part of the small-$q$ coefficient puzzle reflects measurement problems rather than genuinely enormous adjustment costs.

## 9.24 Phase Diagram, Saddle-Path Stability, and the Role of $q$

To study the dynamics of investment more transparently, it is useful to consider a simplified version of the model in which depreciation is set to zero, $\delta = 0$, and adjustment costs depend only on investment, $c(I, k) = c(I)$. The firm's problem then implies two key equations:

$$\dot{k}(t) = I(t),$$

and, from the first-order condition for investment,

$$c'(I(t)) = q(t) - 1.$$

If $c'' > 0$, this condition can be inverted to give

$$I(t) = h(q(t)), \qquad h'(q) > 0.$$

Hence,

$$\dot{k}(t) = h(q(t)).$$

The co-state equation becomes

$$\dot{q}(t) = rq(t) - (1 - \tau)F_k(k(t)).$$

Thus, the dynamic system can be written as

$$\dot{k}(t) = h(q(t)), \qquad \dot{q}(t) = rq(t) - (1 - \tau)F_k(k(t)).$$

**The $\dot{k} = 0$ locus.** Because $I = h(q)$, the capital stock is constant whenever investment is zero. Under the standard normalization $c'(0) = 0$, the investment first-order condition implies

$$I = 0 \iff q = 1.$$

Therefore,

$$\dot{k} = 0 \iff q = 1.$$

This locus is a horizontal line in $(k, q)$-space.

**The $\dot{q} = 0$ locus.** Setting the co-state equation equal to zero yields

$$\dot{q} = 0 \iff q = \frac{1 - \tau}{r} F_k(k).$$

Under diminishing returns, $F_{kk}(k) < 0$, this locus is downward sloping in $(k, q)$-space.

**Steady state.** A steady state $(k^*, q^*)$ satisfies both $\dot{k} = 0$ and $\dot{q} = 0$. Hence,

$$q^* = 1$$

and

$$1 = \frac{1-\tau}{r} F_k(k^*),$$

or equivalently,

$$r = (1-\tau) F_k(k^*).$$

**Intuition.** The condition $q = 1$ means that the shadow value of installed capital is exactly equal to the purchase cost of new capital. At that point the firm has no incentive either to expand or contract its capital stock. The second condition, $r = (1-\tau)F_k(k^*)$, says that in the steady state the after-tax marginal product of capital must equal the required return.

**Direction of motion.** The sign of $\dot{k}$ is governed by $q$:

$$q > 1 \implies I > 0 \implies \dot{k} > 0,$$

$$q < 1 \implies I < 0 \implies \dot{k} < 0.$$

The sign of $\dot{q}$ depends on whether the economy is above or below the $\dot{q} = 0$ locus:

$$q > \frac{1-\tau}{r} F_k(k) \implies \dot{q} > 0,$$

$$q < \frac{1-\tau}{r} F_k(k) \implies \dot{q} < 0.$$

The steady state is a saddle point. There is a unique stable saddle path converging to $(k^*, q^*)$. Since capital is a predetermined state variable, $k$ cannot jump. By contrast, $q$ is a forward-looking co-state variable and may jump immediately. Therefore, for any given initial $k_0$, there is a unique value $q_0$ that places the economy on the stable saddle path.

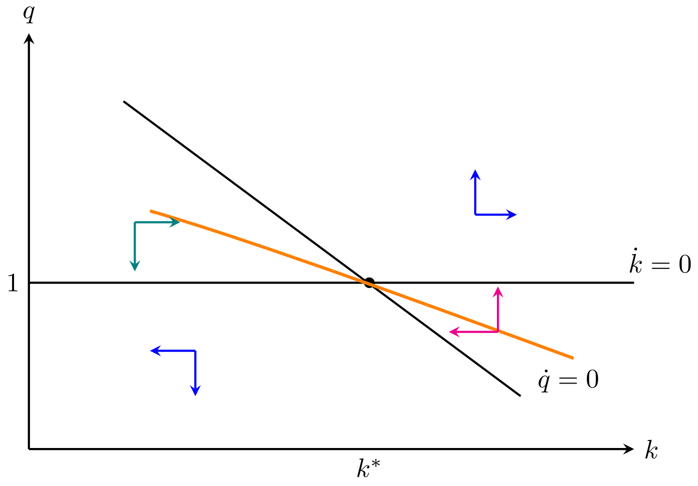

Figure 18: Phase diagram in the $(k, q)$ plane. The horizontal locus is $\dot{k} = 0$, the downward-sloping locus is $\dot{q} = 0$, and the orange curve is the stable saddle path.

**Remark.** This uniqueness is the key economic content of saddle-path stability in this model. The initial capital stock is inherited from the past, while the shadow value $q$ adjusts on impact so that the economy follows the only non-explosive path consistent with optimality.

## 9.25 An Unexpected Permanent Tax Cut

Consider an unexpected permanent decline in the corporate tax rate, so that

$$\Delta\tau < 0.$$

From the $\dot{q} = 0$ locus,

$$q = \frac{1-\tau}{r} F_k(k),$$

a lower tax rate raises the after-tax marginal product of capital for every given $k$. Therefore, the $\dot{q} = 0$ locus shifts upward.

Since the $\dot{k} = 0$ locus remains

$$q = 1,$$

the new steady state is characterized by the same value $q = 1$ but a larger long-run capital stock, say $k^{**} > k^*$. Intuitively, a tax cut raises the after-tax return to capital and therefore increases the firm's desired capital stock.

Because capital is predetermined, the economy cannot jump immediately from $k_0$ to $k^{**}$. Instead, at the moment of the tax cut, $q$ jumps upward so that the economy lands on the new stable saddle path. After that initial jump, capital gradually rises over time toward the new steady state, while $q$ declines back toward its steady-state value of one.

**Intuition.** The tax cut makes installed capital more valuable because future after-tax marginal products are now higher. That higher valuation shows up as an immediate increase in $q$. Once $q > 1$, firms invest more, capital accumulates gradually, and the marginal value of additional capital eventually falls back to $q = 1$ as the economy approaches its new steady state.

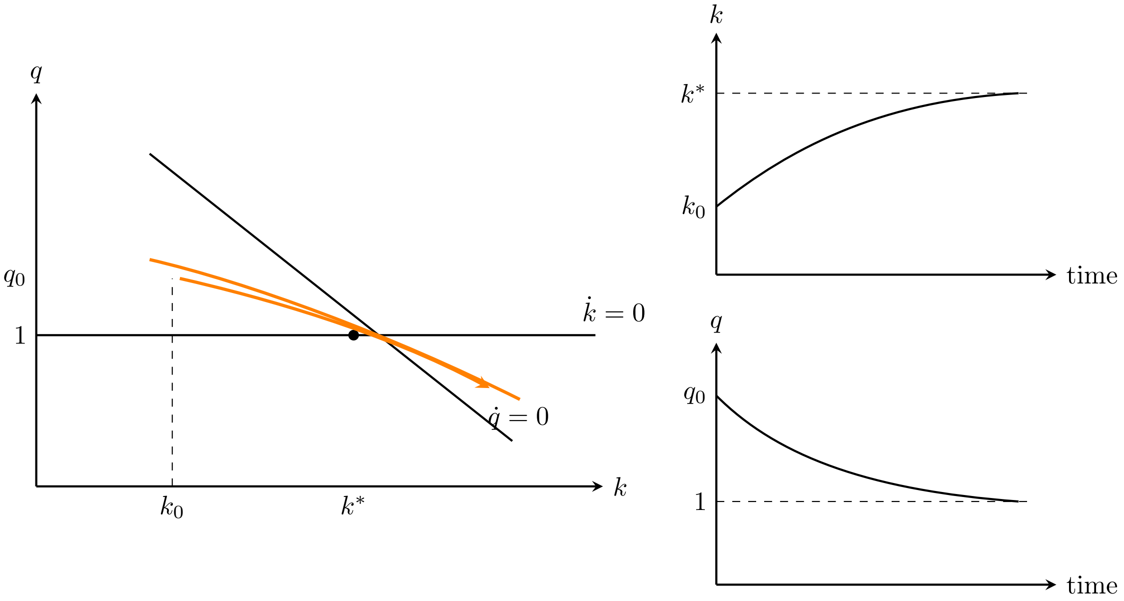

Figure 19: Transition dynamics from an initial capital stock $k_0 < k^*$. Since $k$ is predetermined, it cannot jump. The shadow value $q$ jumps immediately to the saddle path and then declines toward one as capital accumulates.

**Remark.** The phase diagram highlights the distinct roles of the two variables in the model. Capital is predetermined and adjusts only gradually through investment. By contrast, $q$ is forward-looking and can

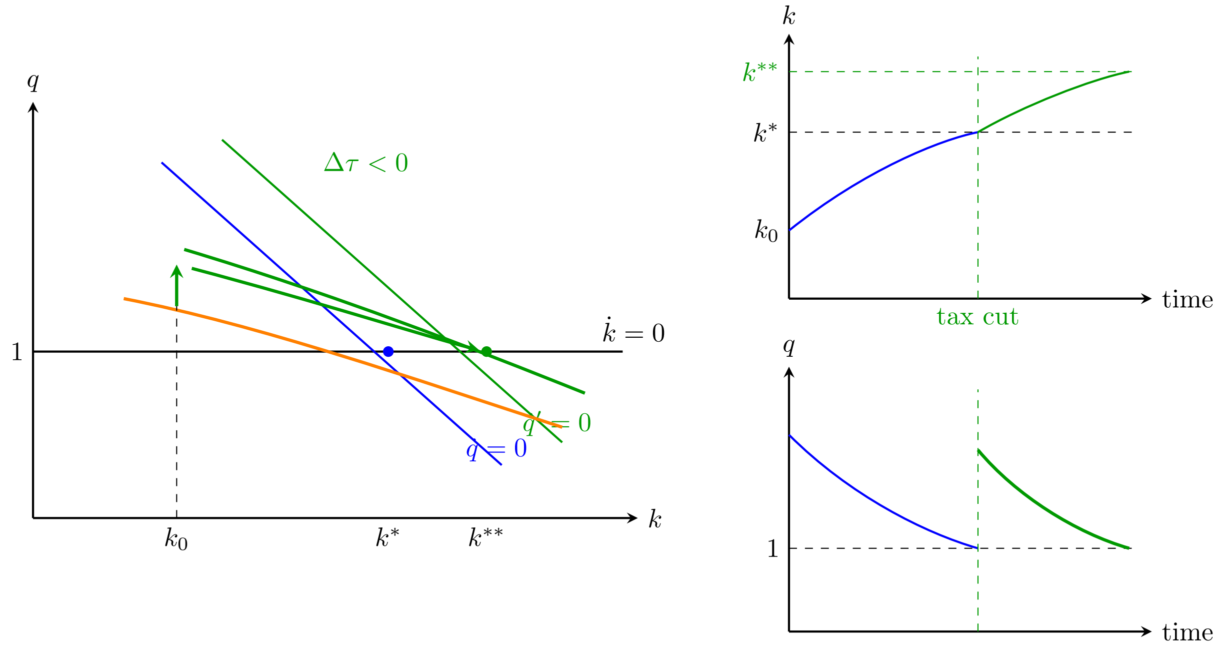

Figure 20: An unexpected permanent tax cut shifts the $\dot{q} = 0$ locus upward, raises the long-run capital stock from $k^*$ to $k^{**}$, and causes an immediate jump in $q$. Capital then adjusts gradually along the new saddle path.

jump immediately in response to news or policy changes. This is why anticipated changes in taxation or profitability are first reflected in asset values and only later in the capital stock itself.

## 9.26 When $q$ Is Not a Sufficient Statistic

The benchmark $q$-theory delivers a simple investment rule when adjustment costs depend only on the investment rate. If

$$c(I_t, K_t) = K_t \phi\left(\frac{I_t}{K_t}\right),$$

then the first-order condition for investment implies

$$q_t = 1 + \phi'\left(\frac{I_t}{K_t}\right),$$

so that, under convexity, one can invert the relation and write

$$\frac{I_t}{K_t} = h(q_t), \qquad h'(q_t) > 0.$$

This clean mapping becomes more delicate once adjustment costs depend on additional shifters. For example, suppose

$$c(I_t, K_t, x_t) = K_t \phi\left(\frac{I_t}{K_t}, x_t\right),$$

where $x_t$ may capture installation wages, energy costs, delivery bottlenecks, or other factors affecting the cost of adjusting capital. Then the first-order condition becomes

$$q_t = 1 + \phi_1\left(\frac{I_t}{K_t}, x_t\right),$$

and therefore

$$\frac{I_t}{K_t} = h(q_t, x_t).$$

In that case, $q_t$ is no longer a sufficient statistic for investment: even holding $q_t$ fixed, changes in $x_t$ may shift the firm's investment demand.

**Remark.** This is an important qualification to the basic empirical specification. Once the adjustment-cost schedule depends on additional state variables or shifters, the tight link between investment and $q$ weakens.

## 9.27 No Adjustment Costs: Marginal Value Versus Marginal Acquisition Cost

The logic of the investment first-order condition is easiest to see in the frictionless case. If there are no adjustment costs,

$$c(I_t, K_t) \equiv 0,$$

then the first-order condition for investment becomes

$$q_t = p_t^k.$$

This condition must be interpreted carefully. It does *not* say that capital has no value. Rather, it says that in the absence of installation frictions, the shadow value of one additional unit of installed capital must equal the marginal acquisition cost of that unit.

More precisely, the relevant equilibrium condition is

$$\text{shadow value of an additional installed unit} = \text{marginal cost of acquiring that unit today.}$$

The object on the left-hand side is $q_t$, the shadow value of installed capital. The object on the right-hand side is the current purchase price of a unit of capital goods, $p_t^k$.

**Intuition.** A useful way to avoid confusion is not to say loosely that "marginal benefit equals marginal cost." The relevant benefit is not gross output today. It is the shadow value $q_t$, namely the increase in firm value generated by one additional unit of installed capital. That shadow value already summarizes the entire future stream of returns generated by the asset. In the frictionless case, optimality requires this shadow value to equal the current acquisition cost $p_t^k$.

**Remark.** If $q_t > p_t^k$, an additional unit of installed capital is worth more than it costs to acquire, so the firm would want to invest more. If $q_t < p_t^k$, capital is worth less than its acquisition cost, so the firm would want less capital. In the frictionless case, this arbitrage logic forces

$$q_t = p_t^k.$$

## 9.28 Valuation of Installed Capital

Without adjustment costs, the shadow value of capital can be written as the present discounted value of the future marginal returns generated by an additional unit of installed capital.

Let $V_t$ denote the marginal value at time $t$ of one additional unit of installed capital. In general, this value is the discounted sum of future after-tax marginal profits generated by that extra unit.

In discrete time, a convenient representation is

$$V_t = \sum_{j=0}^{\infty} \left( \frac{1-\delta}{1+r} \right)^j E_t[(1-\tau_{t+j})\Pi_{K,t+j}] \,,$$

where $\Pi_{K,t+j}$ denotes the marginal contribution of capital to operating profits at date $t+j$.

In continuous time, the corresponding expression is

$$V_t = \int_0^{\infty} e^{-(r+\delta)s}(1-\tau_{t+s})\Pi_K(t+s)\, ds.$$

If operating profits are simply given by production, so that $\Pi(K_t) = F(K_t)$, then $\Pi_K(t) = F_K(K_t)$, and the valuation formula becomes

$$V_t = \int_0^{\infty} e^{-(r+\delta)s}(1-\tau_{t+s})F_K(K_{t+s})\, ds.$$

In the absence of adjustment costs, optimality implies

$$q_t = V_t = p_t^k.$$

Equivalently, the present discounted marginal value of an additional installed unit must equal its replacement cost.

**Remark.** This is the correct non-normalized version of the frictionless condition. The unit of capital is not required to have value one; it is required to have value equal to its acquisition price, $p_t^k$.

## 9.29 Connection with the User-Cost Formula

The valuation condition above is consistent with the user-cost approach developed earlier. In continuous time, the co-state equation for capital is

$$(r+\delta)q_t - \dot{q}_t = (1-\tau_t)\Pi_K(t).$$

In the frictionless case, $q_t = p_t^k$, so substituting into the co-state equation yields

$$(r+\delta)p_t^k - \dot{p}_t^k = (1-\tau_t)\Pi_K(t).$$

If $\Pi_K(t) = p_t^y F_K(K_t)$, this can be written as

$$p_t^y F_K(K_t) = \frac{(r+\delta)p_t^k - \dot{p}_t^k}{1-\tau_t}.$$

Thus, the frictionless $q$-condition and the user-cost formula are simply two ways of expressing the same no-arbitrage logic: the after-tax marginal return from capital must equal its opportunity cost, adjusted for depreciation and capital gains.

## 9.30 Why Depreciation Appears in the Valuation Formula

The appearance of $\delta$ in the valuation formula is sometimes confusing. The key point is that depreciation is not entering as a "benefit." Rather, depreciation reduces the future services delivered by the extra unit of capital installed today.

A unit of capital installed at time $t$ does not remain intact forever. As time passes, only a fraction of that unit survives. Under exponential depreciation, the surviving amount is proportional to $e^{-\delta s}$. Hence the future marginal products generated by today's investment are weighted by both financial discounting and physical decay:

$$e^{-rs}e^{-\delta s} = e^{-(r+\delta)s}.$$

**Intuition.** A higher depreciation rate means that the asset is shorter-lived. Future marginal products therefore receive less weight, because the unit of capital that generates them is gradually disappearing. This is why high-$\delta$ assets have lower marginal value, all else equal.

## 9.31 Investment Tax Subsidies Versus Profit Taxes

It is useful to distinguish two conceptually different tax instruments.

A profit tax affects the *return side* of the valuation condition by scaling down future marginal profits:

$$V_t = \int_0^\infty e^{-(r+\delta)s}(1-\tau_{t+s})\Pi_K(t+s)\,ds.$$

By contrast, an investment tax credit or subsidy affects the *cost side* by lowering the effective acquisition cost of new capital.

If an investment subsidy at rate $s_t$ reduces the effective purchase price of capital from $p_t^k$ to $p_t^k(1-s_t)$, then the frictionless optimality condition becomes

$$q_t = V_t = p_t^k(1-s_t).$$

Thus, profit taxes and investment subsidies enter different sides of the condition. Profit taxes reduce the present discounted value of future returns, whereas investment subsidies reduce the current marginal cost of acquiring capital.

**Remark.** This distinction is conceptually important. A tax on profits changes the payoff from holding installed capital. An investment subsidy changes the price at which new capital can be acquired. They therefore operate through different margins, even if both affect investment incentives.

## 9.32 Long-Lived Capital and the Dynamics of the Capital Stock

The law of motion for capital is

$$K_t = (1-\delta)K_{t-1} + I_t.$$

Iterating forward yields

$$K_t = \sum_{j=0}^{\infty}(1-\delta)^j I_{t-j},$$

so the current capital stock is a weighted sum of past investments, with weights that decay at rate $\delta$.

In continuous time, the same idea can be expressed as

$$K(t) = \int_0^\infty e^{-\delta s} I(t-s)\,ds.$$

**Intuition.** Capital is a stock, not a flow. Current investment adds to the stock, but once installed, that capital remains productive for many future periods. The lower is $\delta$, the more persistent is the contribution of each past investment to the current stock.

This representation is especially useful for understanding long-lived capital goods. When depreciation is low, the current stock reflects the accumulation of many years of past investment, so $K_t$ is large relative to the flow $I_t$.

## 9.33 Why Investment Movements May Have Small Effects on the Capital Stock

Consider a temporary change in current investment, holding past investment fixed. From

$$K_t = (1-\delta)K_{t-1} + I_t,$$

we obtain

$$dK_t = dI_t.$$

Dividing by $K_t$,

$$\frac{dK_t}{K_t} = \frac{dI_t}{K_t} = \frac{I_t}{K_t}\frac{dI_t}{I_t}.$$

In a steady state,

$$K_t = (1-\delta)K_t + I_t \quad \Longrightarrow \quad \frac{I_t}{K_t} = \delta.$$

Hence,

$$\frac{dK_t}{K_t} = \delta\frac{dI_t}{I_t}.$$

This is the key implication: when $\delta$ is small, even large percentage changes in investment translate into small percentage changes in the capital stock.

**Example.** If $\delta = 0.02$, then

$$\frac{dK_t}{K_t} = 0.02\frac{dI_t}{I_t}.$$

Thus, a 10% increase in investment raises the capital stock on impact by only 0.2%.

**Remark.** The correct intuition is therefore not that investment itself moves little when depreciation is low. Rather, because the capital stock is large relative to current investment, fluctuations in investment have only modest short-run effects on the stock of productive capital.

## 9.34 Temporary Investment Tax Credits with Long-Lived Capital

Now consider a temporary investment tax credit. To distinguish it from the profit tax, let $z_t$ denote the subsidy rate on new investment. If the purchase price of capital goods is $p_t^k$, the effective cost of acquiring one new unit is

$$(1-z_t)p_t^k.$$

In the frictionless case, the firm's first-order condition is

$$(1-z_t)p_t^k = V_t,$$

where $V_t$ is the shadow value of an additional installed unit of capital.

Suppose now that capital is long-lived and that the tax credit is temporary. Then the policy mainly affects the timing of current purchases, but it has little effect on the long-run marginal value of installed capital. Over the short horizon of the policy, it is often reasonable to approximate

$$V_t \approx \bar{V},$$

where $\bar{V}$ is approximately constant.

The first-order condition then becomes

$$(1 - z_t)p_t^k \approx \bar{V},$$

or equivalently,

$$p_t^k \approx \frac{\bar{V}}{1 - z_t}.$$

**Intuition.** A temporary subsidy reduces the buyer's net cost of acquiring capital. But if the underlying value of installed capital is roughly unchanged, then buyers are willing to pay a higher gross price. In this sense, part of the subsidy is capitalized into the price of capital goods.

This is the central point for long-lived capital: because the policy is temporary and the asset lasts for many periods, the shadow value $V_t$ does not move much, so the main immediate effect may be a higher purchase price rather than a large increase in the capital stock.

## 9.35 Demand for Investment Goods and Price Capitalization

The previous condition implies that the firm's willingness to pay for new capital is pinned down by

$$p_t^k = \frac{V_t}{1 - z_t}.$$

If $V_t \approx \bar{V}$, investment demand becomes approximately horizontal at the gross price

$$p_t^k = \frac{\bar{V}}{1 - z_t}.$$

Without the subsidy, the corresponding willingness to pay is simply

$$p_t^k = \bar{V}.$$

Thus, a temporary investment tax credit shifts investment demand upward from $\bar{V}$ to $\bar{V}/(1 - z_t)$. If the short-run supply of capital goods is upward sloping, the equilibrium price rises and the equilibrium quantity increases as well. The division of the subsidy between higher prices and higher quantities depends on supply conditions, but the basic force is immediate price capitalization.

**Remark.** The long-lived-capital logic makes current demand for capital goods highly sensitive to the subsidy even when the induced change in the productive capital stock is small. This distinction between the market for investment goods and the dynamics of the installed capital stock is essential.

## 9.36 Connection with Goolsbee (1998)

This logic motivates the empirical approach in Goolsbee (1998), which studies whether investment tax incentives are reflected in quantities, prices, or both. The key empirical idea is to exploit variation across types of capital goods in their exposure to tax incentives.

A stylized specification is

$$\ln I_t^m = \alpha^m + \lambda_t + \beta_I \, ITC_t^m + \varepsilon_t^m,$$

for investment quantities, and

$$\ln p_t^m = \tilde{\alpha}^m + \tilde{\lambda}_t + \beta_P \, ITC_t^m + \nu_t^m,$$

for equipment prices, where $m$ indexes asset types and $ITC_t^m$ measures the relevant investment tax credit.

The central empirical question is whether tax incentives mainly raise the quantity of investment or whether they are substantially capitalized into prices. Empirically, one expects

$$\beta_I > 0,$$

since a larger investment tax credit raises the incentive to purchase the corresponding type of capital good. A positive $\beta_I$ therefore indicates that tax incentives do stimulate investment demand.

At the same time, a central result in this literature is that

$$\beta_P \approx 1.$$

The interpretation is that a large fraction of the tax subsidy is passed through into higher equipment prices rather than lower effective purchase prices for firms. In other words, suppliers capture an important share of the policy through price increases.

This finding is especially natural in the long-lived-capital framework. A temporary tax credit can generate a strong increase in current demand for new capital goods, because firms have an incentive to bring purchases forward in time. But since the induced change in the productive capital stock is small in the short run, much of the immediate adjustment may occur through prices rather than quantities. Thus, $\beta_I > 0$ reflects a positive demand response, while $\beta_P \approx 1$ suggests substantial price capitalization of the subsidy.

**Intuition.** If the government subsidizes the purchase of equipment, the economic incidence of that subsidy need not accrue entirely to investing firms. When supply is upward sloping, part—and potentially most—of the tax benefit may be captured by producers through higher prices. In that case, investment still rises, but the gross price of capital goods rises almost one-for-one with the subsidy.

## 9.37 Hayashi-Type Adjustment Costs and the Prediction of Smooth Investment

With convex adjustment costs, large one-time changes in capital are expensive at the margin. If a firm wants to move its capital stock upward, it is often cheaper to spread the adjustment over time. The model therefore predicts frequent and relatively small investment decisions rather than rare and very large jumps.

This is the basic empirical tension with the Hayashi framework. At the plant level, investment does not appear to be mostly smooth and gradual. Instead, establishments often spend long periods with little adjustment and then undertake large bursts of investment.

**Remark.** A natural question is what happens around $I_t = 0$. If the model allows $I_t < 0$, then firms can actively disinvest by selling or scrapping capital, so small positive and negative adjustments are both possible. In many applications, however, one imposes partial irreversibility or the constraint $I_t \geq 0$, in

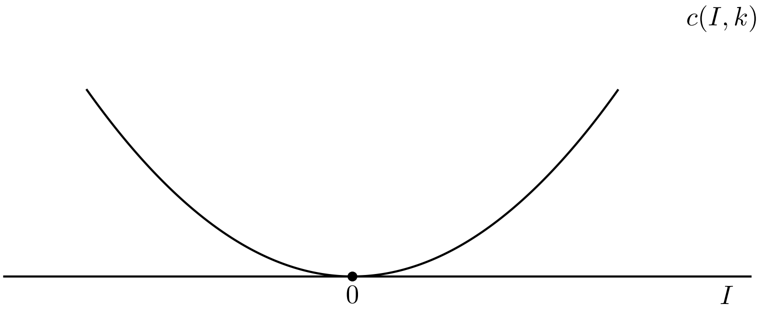

which case firms can choose zero gross investment and let capital fall only through depreciation. This helps rationalize long periods of inaction, but by itself it still does not generate the large positive spikes seen in the data.

## 9.38 Micro Evidence from Plant-Level Data

A central empirical contribution of the literature is to study investment using plant-level data rather than aggregate time series alone. A leading source is the U.S. Census Longitudinal Research Database (LRD), which tracks manufacturing establishments over time.

Doms and Dunne (1998) document that capital adjustment at the plant level is highly uneven over time. Most manufacturing plants experience at least one year in which capital rises by at least 50%. For many establishments, a large share of total investment is concentrated in a small number of years rather than being spread smoothly over time.

This evidence is hard to reconcile with the pure Hayashi model. If adjustment costs were only smooth and convex, one would expect firms to fine-tune capital more gradually. Instead, the data suggest that firms often wait and then adjust in a discrete and concentrated way.

## 9.39 Investment Spikes and Lumpy Investment

To formalize this idea, the literature often defines an *investment spike* as an episode in which plant-level investment is unusually large relative to the existing capital stock. A standard definition is

$$|\frac{K_{i,t}}{K_{i,t-1}}| \geq 0.20.$$

**Remark.** If you want to stay close to the plant-level investment literature, it is better not to use absolute values in the spike definition here. The standard object is a *positive* investment spike, not just a large adjustment of either sign. Large disinvestment episodes can also be studied, but they are conceptually distinct.

Cooper, Haltiwanger, and Power (1999) show that these spikes are not rare events. In any given year, roughly one-fifth of manufacturing plants experience an investment spike. This is why the literature often refers to plant-level capital adjustment as *lumpy investment*: establishments usually adjust little, but occasionally undertake very large capital expenditures.

**Intuition.** The key idea is that firms do not necessarily move capital smoothly in every period. Instead, they may tolerate being away from their frictionless optimum for some time and then adjust in a large step once the gains from adjustment are high enough.

## 9.40 From Plant-Level Lumpiness to Aggregate Investment

An important question is whether these plant-level spikes matter for aggregate fluctuations. Gourio and Kashyap (2007) show that they do.

A useful decomposition distinguishes between:

- the *intensive margin*: the average size of investment among plants that are adjusting;
- the *extensive margin*: the fraction or number of plants that undertake an investment spike in a given period.

Their main finding is that aggregate investment fluctuations are driven primarily by the extensive margin. That is, aggregate investment rises mainly because more plants undertake spikes, not because the average spike among adjusters becomes much larger.

**Intuition.** This is analogous to other state-dependent adjustment problems in macroeconomics. Aggregate movements need not come from every unit doing a little more. They can instead arise because many units cross an adjustment threshold at roughly the same time.

Thus, the plant-level evidence can be summarized as follows: investment is often low or near zero for long stretches, but from time to time plants undertake large spikes, and aggregate investment co-moves strongly with the number of plants experiencing those spikes.

## 9.41 Implication for Theory

Taken together, these findings suggest that pure Hayashi-style convex adjustment costs are not sufficient to describe plant-level investment dynamics. The micro evidence points instead toward models with non-convexities, fixed costs, partial irreversibility, or generalized $(S, s)$-type adjustment rules.

The broad lesson is that the distinction between smooth adjustment and lumpy adjustment is not a minor quantitative detail. It changes both the microeconomic description of firm behavior and the way one thinks about aggregate investment fluctuations.

## 9.42 Kinked Adjustment Costs and Partial Irreversibility

We now consider a different form of adjustment costs: instead of smooth convex costs, suppose there is a wedge between the price at which the firm buys capital and the price at which it can sell it.

Let $p^k$ denote the purchase price of capital. If the firm buys one unit of capital, it pays $p^k$. If instead it sells capital, it only recovers a fraction $\alpha p^k$, where

$$0 \leq \alpha < 1.$$

Thus, capital is partially irreversible: selling capital is possible, but the firm does not recover the full purchase price.

**Simplifying assumptions.** To isolate the logic, suppose:

1. the purchase price is normalized to one,
$$p^k = 1;$$
2. there is no depreciation,
$$\delta = 0;$$

3. current profits are given by $\pi(k, \varepsilon)$, with

$$\pi_k(k,\varepsilon) > 0, \qquad \pi_{kk}(k,\varepsilon) < 0,$$

$$\pi_\varepsilon(k,\varepsilon) > 0, \qquad \pi_{k\varepsilon}(k,\varepsilon) \geq 0;$$

4. the shock $\varepsilon$ is i.i.d.;

5. newly installed capital becomes productive immediately.

**Intuition.** The assumptions on $\pi$ have a natural economic interpretation. More capital raises profits, but at a diminishing rate. A higher productivity shock raises profits directly, and the condition $\pi_{k\varepsilon} \geq 0$ says that capital is especially valuable in good states: when productivity is high, the marginal payoff from an extra unit of capital is also high.

Because $\delta = 0$, net investment is simply

$$I = k' - k.$$

The adjustment-cost schedule is therefore

$$\phi(k'-k) = \begin{cases} k' - k, & \text{if } k' \geq k, \\ \alpha(k'-k), & \text{if } k' < k. \end{cases}$$

When the firm expands, it pays one unit per unit of capital. When it contracts, it receives only $\alpha < 1$ per unit sold.

**Remark.** This is why the adjustment-cost function is *kinked* at zero. The marginal cost of increasing capital is 1, while the marginal value of reducing capital is only $\alpha$. Equivalently, buying and selling capital take place at different prices.

Let the state variables be $(k, \varepsilon)$. Since newly chosen capital is productive immediately, the Bellman equation is

$$V(k,\varepsilon) = \max_{k'} \left\{ \pi(k',\varepsilon) - \phi(k'-k) + \beta \mathbb{E}\big[V(k',\varepsilon')\big] \right\}.$$

Because $\varepsilon$ is i.i.d., it is convenient to define

$$v(k) \equiv \mathbb{E}_{\varepsilon'}[V(k,\varepsilon')] = \int V(k,\varepsilon') f(\varepsilon')\, d\varepsilon'.$$

Then the Bellman equation becomes

$$V(k,\varepsilon) = \max_{k'} \left\{ \pi(k',\varepsilon) - \phi(k'-k) + \beta v(k') \right\}.$$

**Step 1: first-order conditions in the buying and selling regions.** Suppose first that it is optimal to *buy* capital, so that $k' > k$. In that region,

$$\phi(k'-k) = k' - k,$$

and the first-order condition is

$$\pi_k(k',\varepsilon) - 1 + \beta v_k(k') = 0.$$

Now suppose instead that it is optimal to *sell* capital, so that $k' < k$. In that region,

$$\phi(k' - k) = \alpha(k' - k),$$

and the first-order condition is

$$\pi_k(k', \varepsilon) - \alpha + \beta v_k(k') = 0.$$

Define

$$m(k', \varepsilon) \equiv \pi_k(k', \varepsilon) + \beta v_k(k').$$

Then the two first-order conditions can be written compactly as

$$m(k', \varepsilon) = 1 \qquad \text{if } k' > k,$$

and

$$m(k', \varepsilon) = \alpha \qquad \text{if } k' < k.$$

**Step 2: why an inaction region appears.** Since profits are concave and $v$ is also concave, $m(k', \varepsilon)$ is decreasing in $k'$. Therefore there exist two thresholds, $k_L(\varepsilon)$ and $k_H(\varepsilon)$, defined by

$$m(k_L(\varepsilon), \varepsilon) = 1, \qquad m(k_H(\varepsilon), \varepsilon) = \alpha.$$

Because $1 > \alpha$ and $m$ is decreasing, we have

$$k_L(\varepsilon) \leq k_H(\varepsilon).$$

This immediately implies a band of inaction.

If the current capital stock is very low, $k < k_L(\varepsilon)$, the firm wants to buy capital and adjusts up to the lower trigger $k_L(\varepsilon)$.

If the current capital stock is very high, $k > k_H(\varepsilon)$, the firm wants to sell capital and adjusts down to the upper trigger $k_H(\varepsilon)$.

If instead

$$k_L(\varepsilon) \leq k \leq k_H(\varepsilon),$$

then it is optimal not to adjust at all and simply choose

$$k' = k.$$

**Intuition.** The wedge between buying and selling prices creates a region in which the firm tolerates being away from the frictionless optimum. If capital is slightly too low, buying is not worth paying the full purchase price. If capital is slightly too high, selling is not attractive because the firm only recovers $\alpha < 1$. Hence the firm remains inactive over an interval of states.

**Step 3: policy rule.** The optimal policy can therefore be written as

$$k'(k, \varepsilon) = \begin{cases} k_L(\varepsilon), & \text{if } k < k_L(\varepsilon), \\ k, & \text{if } k_L(\varepsilon) \leq k \leq k_H(\varepsilon), \\ k_H(\varepsilon), & \text{if } k > k_H(\varepsilon). \end{cases}$$

Thus, kinked adjustment costs imply a *state-dependent* policy with an *inaction region.* Small deviations from the desired capital stock are not corrected immediately. Adjustment occurs only when the state is sufficiently far from the no-adjustment band.

**Remark.** This is the simplest version of an $(S, s)$-type logic for capital adjustment. The firm does not continuously fine-tune its capital stock. Instead, it waits while the state remains inside an inaction region, and adjusts only when the gap becomes large enough.

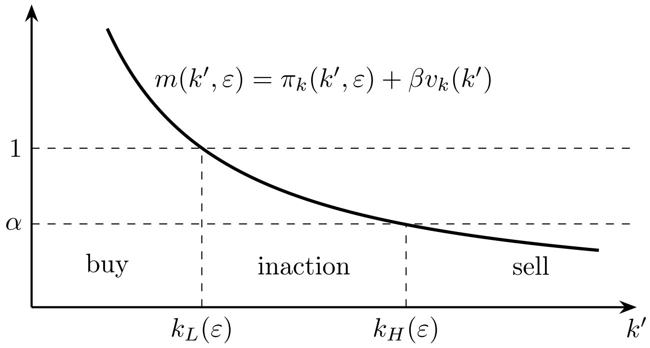

Figure 21: Kinked adjustment costs and the inaction region.

**Comparative statics with respect to productivity.** Suppose $\hat{\varepsilon} > \varepsilon$. Since $\pi_{k\varepsilon} \geq 0$, a higher productivity state shifts

$$m(k', \varepsilon) \equiv \pi_k(k', \varepsilon) + \beta v_k(k')$$

upward for every $k'$. Therefore both adjustment triggers move to the right:

$$k_L(\hat{\varepsilon}) > k_L(\varepsilon), \qquad k_H(\hat{\varepsilon}) > k_H(\varepsilon).$$

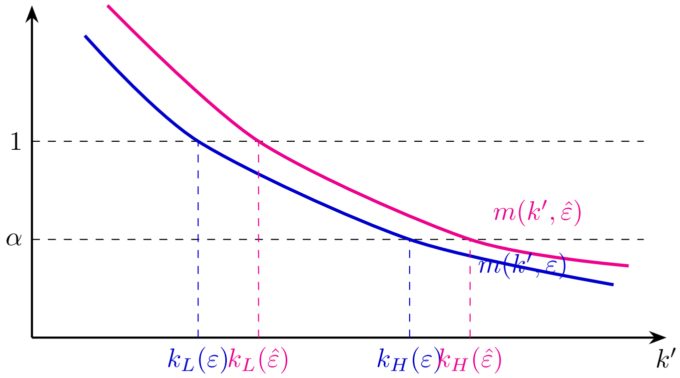

Figure 22: A higher productivity shock shifts $m(k', \varepsilon)$ upward and moves both adjustment triggers to the right.

**Policy function and value function.** The policy rule is

$$k'(k, \varepsilon) = \begin{cases} k_L(\varepsilon), & \text{if } k < k_L(\varepsilon), \\ k, & \text{if } k_L(\varepsilon) \leq k \leq k_H(\varepsilon), \\ k_H(\varepsilon), & \text{if } k > k_H(\varepsilon). \end{cases}$$

Using this policy, the value function can be written piecewise as

$$V(k,\varepsilon) = \begin{cases} \pi\big(k_L(\varepsilon),\varepsilon\big) - k_L(\varepsilon) + k + \beta v\big(k_L(\varepsilon)\big), & \text{if } k < k_L(\varepsilon), \\ \pi(k,\varepsilon) + \beta v(k), & \text{if } k_L(\varepsilon) \le k \le k_H(\varepsilon), \\ \pi\big(k_H(\varepsilon),\varepsilon\big) - \alpha k_H(\varepsilon) + \alpha k + \beta v\big(k_H(\varepsilon)\big), & \text{if } k > k_H(\varepsilon). \end{cases}$$

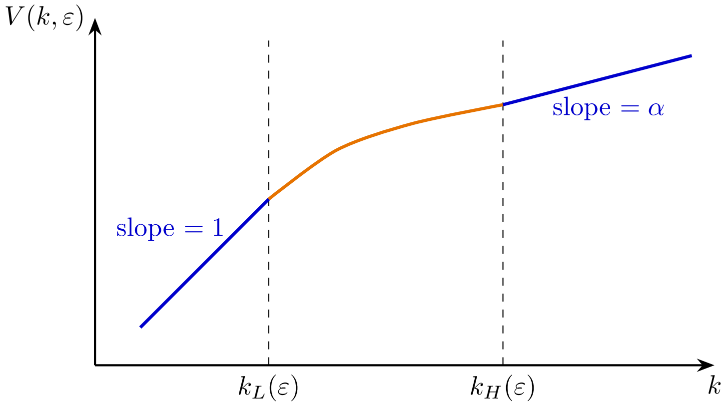

Figure 23: The value function is linear in the adjustment regions and curved in the inaction region.

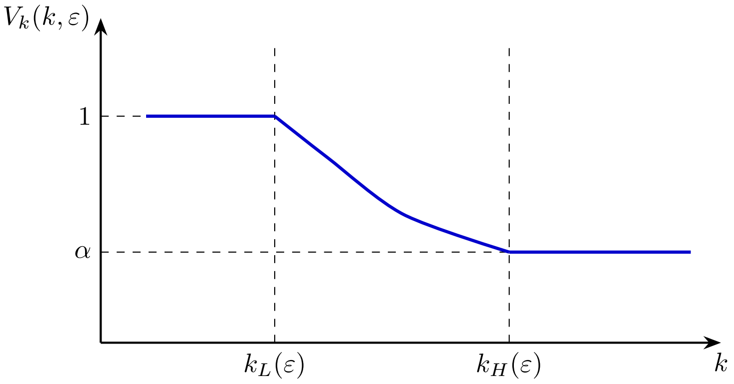

Figure 24: The derivative $V_k(k,\varepsilon)$ is continuous but its slope changes at the triggers. Hence $V$ is generally $C^1$, but not $C^2$.

**Connection with $q$-theory.** This setup has a natural interpretation in terms of marginal $q$. Define

$$q(k,\varepsilon) \equiv V_k(k,\varepsilon),$$

that is, the shadow value of one more unit of installed capital.

Using the piecewise form of the value function, we obtain

$$q(k,\varepsilon) = V_k(k,\varepsilon) = \begin{cases} 1, & \text{if } k < k_L(\varepsilon), \\ \pi_k(k,\varepsilon) + \beta v_k(k), & \text{if } k_L(\varepsilon) \le k \le k_H(\varepsilon), \\ \alpha, & \text{if } k > k_H(\varepsilon). \end{cases}$$

Thus, marginal $q$ is pinned down by the transaction prices in the adjustment regions: when the firm is buying capital, the shadow value of installed capital is 1, the purchase price; when the firm is selling capital, the shadow value is $\alpha$, the resale price.

Inside the inaction region,

$$\alpha < q(k, \varepsilon) < 1,$$

so the shadow value lies strictly between the purchase and resale prices. This is exactly why the firm does not adjust: the marginal value of installed capital is too low to justify buying, but too high to justify selling.

At the trigger points, smooth pasting implies

$$V_k\big(k_L(\varepsilon), \varepsilon\big) = 1, \qquad V_k\big(k_H(\varepsilon), \varepsilon\big) = \alpha.$$

**Remark.** In the standard convex-adjustment-cost version of $q$-theory, $q$ moves smoothly with investment through a first-order condition. Here, by contrast, partial irreversibility creates a band: $q = 1$ in the buying region, $q = \alpha$ in the selling region, and $q \in (\alpha, 1)$ in the no-adjustment region.

## 9.43 Fixed Adjustment Costs

We now consider another nonconvex adjustment technology: a fixed adjustment cost. Suppose the firm must pay a fixed amount

$$F > 0$$

whenever it changes its capital stock, that is, whenever $k' \neq k$.

**Assumptions.** As before, assume:

1. no depreciation, so $\delta = 0$;
2. newly installed capital becomes productive immediately;
3. the purchase price of capital is normalized to one;
4. the resale price is also one, so there is no partial irreversibility here;
5. current profits are $\pi(k, \varepsilon)$, increasing and concave in $k$.

The state is $(k, \varepsilon)$, and the Bellman equation can be written as

$$V(k, \varepsilon) = \max\left\{\pi(k, \varepsilon) + \beta v(k),\ \max_{k'}\Big[\pi(k', \varepsilon) - (k' - k) - F + \beta v(k')\Big]\right\}.$$

The first term corresponds to *inaction*: the firm keeps $k' = k$ and avoids paying the fixed cost. The second term corresponds to *adjustment*: the firm pays $F$ and chooses a new capital stock $k'$.

**Adjustment target.** Conditional on adjusting, the firm's problem is

$$\max_{k'}\Big\{\pi(k', \varepsilon) - k' + \beta v(k')\Big\} + k - F.$$

Since the term $k - F$ does not depend on $k'$, the optimal choice under adjustment solves

$$k^*(\varepsilon) \in \arg\max_{k'}\Big\{\pi(k', \varepsilon) - k' + \beta v(k')\Big\}.$$

Under concavity, this target is unique and satisfies

$$\pi_k(k^*(\varepsilon), \varepsilon) + \beta v_k(k^*(\varepsilon)) = 1.$$

Define

$$A(\varepsilon) \equiv \pi(k^*(\varepsilon), \varepsilon) - k^*(\varepsilon) - F + \beta v(k^*(\varepsilon)).$$

Then the value of adjusting can be written as

$$V^A(k, \varepsilon) = A(\varepsilon) + k.$$

**Remark.** This is the key implication of a fixed adjustment cost. Once the firm decides to pay $F$, the optimal target $k^*(\varepsilon)$ does not depend on current capital $k$. Hence adjustment is *lumpy*: the firm jumps directly to $k^*(\varepsilon)$.

The value of inaction is

$$V^N(k, \varepsilon) = \pi(k, \varepsilon) + \beta v(k).$$

Therefore,

$$V(k, \varepsilon) = \max\bigl\{V^N(k, \varepsilon),\, V^A(k, \varepsilon)\bigr\} = \max\left\{\pi(k, \varepsilon) + \beta v(k),\ A(\varepsilon) + k\right\}.$$

**Inaction region.** Since $V^N(k, \varepsilon)$ is concave in $k$, while $V^A(k, \varepsilon) = A(\varepsilon) + k$ is linear with slope 1, there may exist two thresholds

$$k_L(\varepsilon) < k_H(\varepsilon)$$

such that

$$\pi(k_j(\varepsilon), \varepsilon) + \beta v(k_j(\varepsilon)) = A(\varepsilon) + k_j(\varepsilon), \qquad j \in \{L, H\},$$

and, by tangency,

$$\pi_k(k_j(\varepsilon), \varepsilon) + \beta v_k(k_j(\varepsilon)) = 1, \qquad j \in \{L, H\}.$$

The optimal policy is then

$$k'(k, \varepsilon) = \begin{cases} k^*(\varepsilon), & \text{if } k < k_L(\varepsilon), \\ k, & \text{if } k_L(\varepsilon) \leq k \leq k_H(\varepsilon), \\ k^*(\varepsilon), & \text{if } k > k_H(\varepsilon). \end{cases}$$

**Intuition.** With a fixed adjustment cost, small deviations from the target are not worth correcting. But once the firm adjusts, it is optimal to jump all the way to the same target $k^*(\varepsilon)$. This is why the policy exhibits *inaction* together with *lumpy adjustment.*

**Why is the intercept $A(\varepsilon)$ positive?** The vertical intercept of the adjustment branch is

$$V^A(0, \varepsilon) = A(\varepsilon).$$

Thus, $A(\varepsilon)$ is positive if and only if adjusting from zero capital yields a positive value:

$$A(\varepsilon) > 0 \quad \Longleftrightarrow \quad \max_{k'}\Bigl\{\pi(k', \varepsilon) - k' + \beta v(k')\Bigr\} > F.$$

Equivalently, if the left threshold satisfies $k_L(\varepsilon) > 0$, then $k = 0$ lies in the adjustment region, so adjustment must dominate inaction at $k = 0$:

$$A(\varepsilon) = V^A(0,\varepsilon) > V^N(0,\varepsilon) = \pi(0,\varepsilon) + \beta v(0).$$

Under the common normalization $\pi(0,\varepsilon) = 0$ and $v(0) \geq 0$, this implies

$$A(\varepsilon) > 0.$$

**Remark.** So the positive intercept in the figure is not automatic for every parameter configuration. It holds in the states for which adjustment is optimal even from very low capital levels.

**Where does the fixed cost appear in the graph?** At the adjustment target $k^*(\varepsilon)$,

$$V^N(k^*(\varepsilon),\varepsilon) = \pi(k^*(\varepsilon),\varepsilon) + \beta v(k^*(\varepsilon)),$$

while

$$V^A(k^*(\varepsilon),\varepsilon) = A(\varepsilon) + k^*(\varepsilon) = \pi(k^*(\varepsilon),\varepsilon) + \beta v(k^*(\varepsilon)) - F.$$

Hence

$$V^N(k^*(\varepsilon),\varepsilon) - V^A(k^*(\varepsilon),\varepsilon) = F.$$

So the fixed cost is exactly the vertical gap between the no-adjustment branch and the adjustment branch at $k^*(\varepsilon)$.

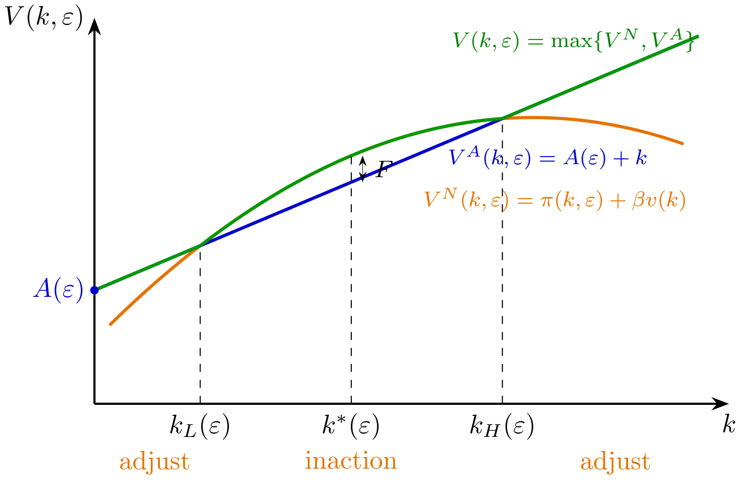

Figure 25: With fixed adjustment costs, the firm either remains inactive or pays the fixed cost and jumps to the target $k^*(\varepsilon)$. The adjustment branch is linear in current capital, and the vertical gap at $k^*(\varepsilon)$ is exactly the fixed cost $F$.

## 9.44 Extension: Persistent Shocks, Depreciation, and the Cross-Section of Firms

We now extend the fixed-adjustment-cost model to a setting in which profitability shocks are persistent rather than i.i.d. The key economic point is that persistence changes the continuation value of capital. A high current shock is likely to remain high for some time, so being close to the desired capital stock becomes more valuable. Likewise, a low current shock is likely to remain low, so excess capital becomes more costly for more than one period.

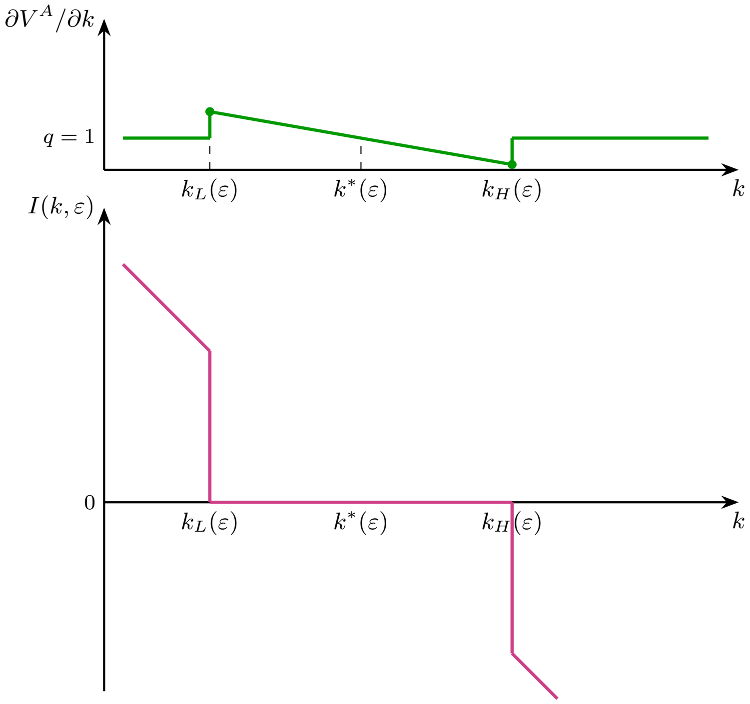

Figure 26: Lower panels associated with the fixed-adjustment-cost problem. The upper panel plots the slope of the value function. In the inaction region, the firm follows the no-adjustment branch, while in the adjustment regions the slope equals that of the adjustment branch. The lower panel plots the investment rule when $\delta = 0$: the firm does not invest inside the inaction region and, whenever it adjusts, chooses investment to jump directly to the target capital stock $k^*(\varepsilon)$.

This changes both the geometry of the value function and the firm's adjustment policy. In the i.i.d. benchmark, the decision to adjust is relatively myopic: the firm compares today's benefit from adjusting with today's fixed cost. With persistent shocks, the same comparison becomes much more forward-looking, because today's capital position affects profits over many future periods.

### 9.44.1 A continuous-time formulation

Let $\varepsilon_t$ denote a persistent profitability shock evolving as a diffusion:

$$d\varepsilon_t = \mu(\varepsilon_t)\,dt + \sigma(\varepsilon_t)\,dW_t.$$

When the firm does not adjust, capital depreciates at rate $\delta$:

$$\dot{k}_t = -\delta k_t.$$

Let $V(k,\varepsilon)$ be the value function. Inside the inaction region, the Hamilton–Jacobi–Bellman equation is

$$rV^N(k,\varepsilon) = \pi(k,\varepsilon) - \delta k\, V_k^N(k,\varepsilon) + \mu(\varepsilon)V_\varepsilon^N(k,\varepsilon) + \frac{1}{2}\sigma(\varepsilon)^2 V_{\varepsilon\varepsilon}^N(k,\varepsilon).$$

If the firm adjusts, it pays a fixed cost $F$ and chooses a new capital stock $k'$. Normalizing the purchase/sale price of capital to one, the adjustment branch is

$$V^A(k,\varepsilon) = \max_{k'} \left\{ V^N(k',\varepsilon) - F - (k'-k) \right\}.$$

Rearranging,

$$V^A(k,\varepsilon) = k + \max_{k'} \left\{ V^N(k',\varepsilon) - k' \right\} - F.$$

Thus, conditional on adjustment, the value function is linear in current capital:

$$V^A(k,\varepsilon) = A(\varepsilon) + k,$$

for some function $A(\varepsilon)$. The overall value function is therefore

$$V(k,\varepsilon) = \max \left\{ V^N(k,\varepsilon),\, V^A(k,\varepsilon) \right\}.$$

**Intuition.** The adjustment branch is linear in $k$ because once the firm has decided to pay the fixed cost, the current capital stock only matters through how much net investment must be purchased or sold. By contrast, the no-adjustment branch embeds the full dynamic continuation value.

### 9.44.2 Why persistence changes the geometry

With i.i.d. shocks, the value of waiting is limited, because tomorrow's shock is unrelated to today's shock. With persistent shocks, waiting has a richer option value: a firm that is close to a trigger may optimally postpone adjustment because the shock may mean-revert, or because current favorable conditions may last. For this reason, persistent shocks tend to smooth the value function, especially near the boundaries of the inaction region.

Under a diffusion specification, one typically obtains both value matching and smooth pasting at the adjustment thresholds:

$$V^N(k^L(\varepsilon),\varepsilon) = V^A(k^L(\varepsilon),\varepsilon), \qquad V_k^N(k^L(\varepsilon),\varepsilon) = V_k^A(k^L(\varepsilon),\varepsilon),$$

and similarly at $k^H(\varepsilon)$. Relative to the i.i.d. benchmark, the kinks become less sharp.

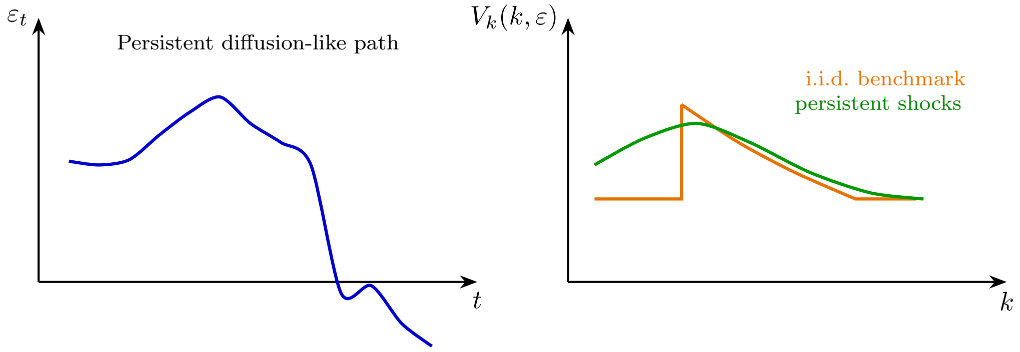

Figure 27: Persistence changes the dynamic problem. The left panel illustrates a persistent shock path. The right panel gives a qualitative comparison of the derivative of the value function: with persistent shocks, the continuation value smooths the geometry near the adjustment thresholds.

### 9.44.3 Homogeneity and the normalized state variable

A useful way to organize the problem is to exploit proportionality between the optimal capital stock and the profitability shock. Suppose the frictionless static target is proportional to $\varepsilon$:

$$k^s(\varepsilon) = \alpha\varepsilon.$$

Likewise, suppose the inaction thresholds are proportional to the same shock:

$$k^H(\varepsilon) = k^s(\varepsilon)(1+\xi) = \alpha(1+\xi)\varepsilon,$$

$$k^L(\varepsilon) = k^s(\varepsilon)(1-\xi) = \alpha(1-\xi)\varepsilon.$$

Now define the normalized state variable

$$z \equiv \frac{k}{\varepsilon}.$$

Then the target and thresholds become constants in $z$-space:

$$z^s \equiv \frac{k^s(\varepsilon)}{\varepsilon} = \alpha, \qquad z^H \equiv \frac{k^H(\varepsilon)}{\varepsilon} = \alpha(1+\xi), \qquad z^L \equiv \frac{k^L(\varepsilon)}{\varepsilon} = \alpha(1-\xi).$$

This is a major simplification. Instead of tracking a target $k^s(\varepsilon)$ that moves with the shock, we track the ratio $z = k/\varepsilon$, and the problem becomes an $(s, S)$-type problem with constant boundaries:

$$z \in [z^L, z^H] \quad \Longrightarrow \quad \text{inaction},$$

$$z < z^L \quad \Longrightarrow \quad \text{upward adjustment},$$

$$z > z^H \quad \Longrightarrow \quad \text{downward adjustment}.$$

**Intuition.** The firm does not care only about capital in levels. What matters is whether capital is high or low relative to the current profitability state. The variable $z = k/\varepsilon$ measures exactly that mismatch.

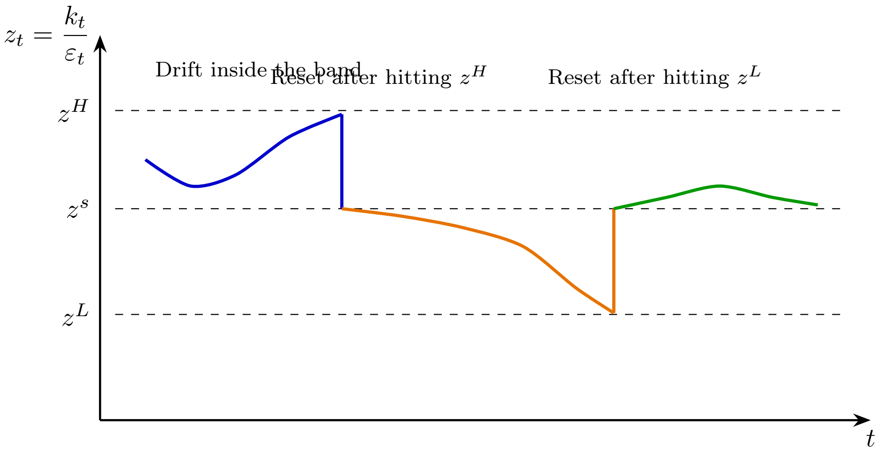

Figure 28: Dynamics of the normalized state $z_t = k_t/\varepsilon_t$. The firm drifts inside the inaction band $[z^L, z^H]$. When a trigger is hit, the firm pays the fixed cost and resets toward the target $z^s$.

### 9.44.4 What depreciation changes

Depreciation adds a mechanical force pushing capital downward between adjustments. In discrete time, if the firm does not adjust,

$$k_{t+1} = (1-\delta)k_t,$$

so in logs

$$\log k_{t+1} = \log k_t + \log(1-\delta).$$

Thus, even if the profitability shock were temporarily constant, non-adjusting firms would drift leftward in $\log k$-space.

In terms of the normalized state $z_t = k_t/\varepsilon_t$, depreciation tends to push the firm toward the lower trigger $z^L$, unless a sufficiently adverse movement in $\varepsilon_t$ pushes it upward instead. Therefore, with $\delta > 0$, even a firm facing a constant profitability environment would not remain forever at the same point in the inaction region.

**Intuition.** Without depreciation, the capital stock can remain fixed until the shock moves enough to trigger adjustment. With depreciation, time itself moves the state. The firm may be forced to reinvest even without a large new shock.

### 9.44.5 Why fixed costs generate overshooting

Let the static flow objective be

$$\Pi(k, \varepsilon) \equiv \pi(k, \varepsilon) - (r + \delta)k.$$

Its frictionless maximizer $k^s(\varepsilon)$ solves

$$\Pi_k(k^s(\varepsilon), \varepsilon) = 0, \qquad \text{equivalently} \qquad \pi_k(k^s(\varepsilon), \varepsilon) = r + \delta.$$

If adjustment were costless, the firm would always set $k = k^s(\varepsilon)$. With a fixed cost $F$, however, this is no longer generally optimal. Once the firm pays the fixed cost, it may prefer to jump beyond the static optimum in order to remain inactive for longer afterward.

We therefore distinguish between:

$$k^s(\varepsilon) \quad \text{(static frictionless optimum)}$$

and

$$k^R(\varepsilon) \quad \text{(reset point after paying the fixed cost).}$$

For upward adjustment, one typically has

$$k^R(\varepsilon) > k^s(\varepsilon)$$

when depreciation is important. The firm intentionally overshoots because capital will gradually erode after the reset, and it is optimal to spread the fixed adjustment cost over a longer inaction spell.

**Intuition.** The static optimum tells us where the firm would like to be *today.* The reset point tells us where the firm wants to land *given that it is costly to readjust again tomorrow.* The latter can lie above the former.

### 9.44.6 Volatility and the option value of waiting

An increase in the volatility of shocks tends to increase the option value of waiting. Suppose a firm is close to the upper trigger $z^H$. If volatility is low, the firm expects little chance of a quick reversal and may prefer to adjust soon. If volatility is high, the firm may reasonably hope that a favorable shock realization will move it back toward the interior of the inaction region without paying the fixed cost.

For this reason, a standard qualitative comparative static in these models is:

$$\uparrow \sigma_\varepsilon \quad \Longrightarrow \quad \text{higher option value of waiting} \quad \Longrightarrow \quad \text{wider inaction region.}$$

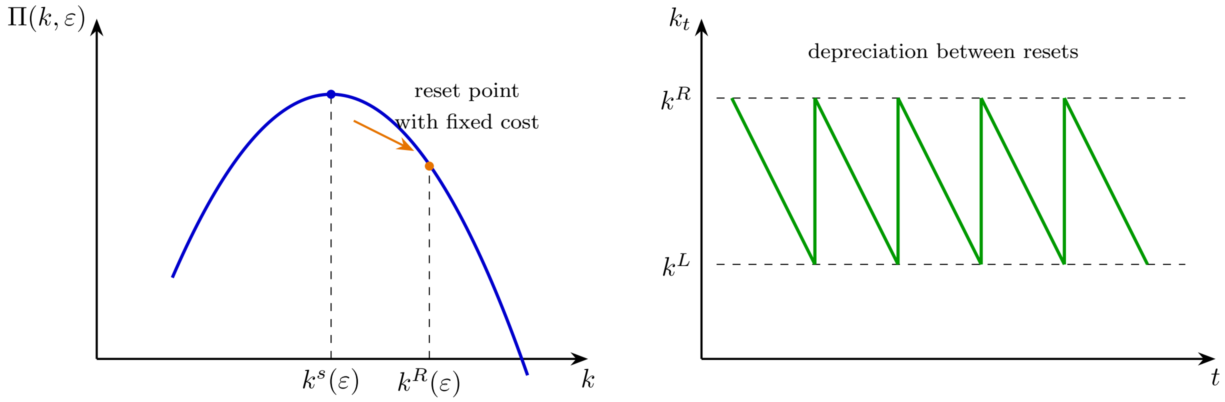

Figure 29: Left: the reset point after paying the fixed cost may exceed the static frictionless optimum. Right: with depreciation, the firm's capital stock gradually erodes until it reaches the lower trigger, at which point it jumps back to the reset point.

**Intuition.** Volatility creates a real option. If the environment may improve on its own, immediate intervention becomes less attractive. The firm is more willing to tolerate temporary misalignment.

### 9.44.7 A continuum of firms and the cross-sectional distribution

The previous discussion describes the optimal policy of an individual firm. For aggregate investment, however, we need the distribution of firms across states. This is where the lecture's final point becomes crucial: the cross-sectional distribution of capital becomes a state variable.

To see the logic, imagine a continuum of firms that have recently adjusted and are then drifting through the inaction region under depreciation. In discrete time, non-adjusters satisfy

$$k_{t+1} = (1-\delta)k_t, \qquad \log k_{t+1} = \log k_t + \log(1-\delta).$$

Hence, in $\log k$-space, non-adjusting firms move leftward at constant speed. If firms are continuously being reset near the upper end of the inaction region and then drifting downward until the lower trigger, the stationary cross-sectional distribution in $\log k$-space is naturally close to uniform.

Let $g_t(\log k)$ denote the density of firms over $\log k$. Aggregate investment is then

$$I_t^{\text{agg}} = \int i(k,\varepsilon_t)\, g_t(\log k)\, d\log k.$$

This expression makes clear that aggregate investment depends not only on the average capital stock, but also on how many firms are close to an adjustment threshold.

**Intuition.** A representative-firm summary is insufficient. Two economies with the same aggregate capital may react very differently to the same policy if one has many firms near the investment trigger and the other does not.

### 9.44.8 Temporary investment tax credits and why the distribution matters

Now consider a temporary investment tax credit (ITC), say at rate $\tau_t > 0$, which temporarily lowers the effective price of new investment. For a firm that undertakes an upward adjustment from $k$ to $k'$, the effective intervention cost becomes

$$(1-\tau_t)(k'-k) + F.$$

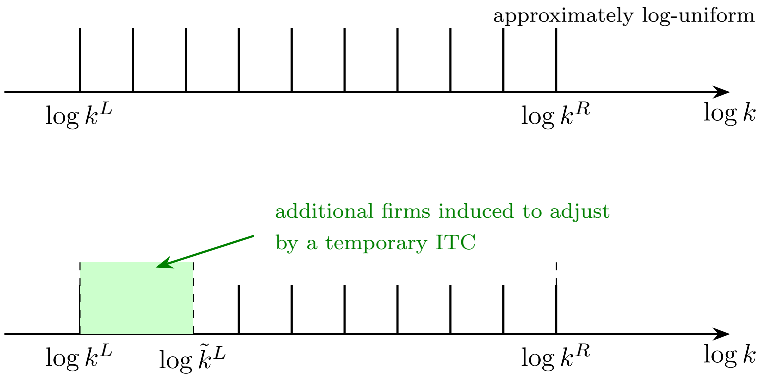

Figure 30: Top: if firms drift through the inaction region under depreciation, the stationary cross-sectional density in $\log k$-space is naturally close to uniform. Bottom: a temporary investment tax credit can move the effective adjustment trigger inward, causing additional firms near the lower boundary to adjust immediately.

This temporary subsidy makes adjustment attractive for some firms that would otherwise have remained inactive. In terms of trigger policies, the lower investment threshold moves inward:

$$k^L \quad \longrightarrow \quad \tilde{k}^L, \qquad \tilde{k}^L > k^L,$$

so that firms with capital stocks in the interval

$$k \in [k^L, \tilde{k}^L]$$

are newly induced to invest.

The aggregate effect depends on the mass of firms in that interval:

$$I_t^{\text{agg,extra}} = \int_{\log k^L}^{\log \tilde{k}^L} \left[k^R(\varepsilon_t) - k\right] g_t(\log k)\, d\log k.$$

This is the precise sense in which the distribution of capital becomes a state variable. The response to policy depends on the entire cross section, not merely on averages.

**Intuition.** A temporary ITC does not move all firms equally. It mainly affects *marginal firms*, namely those already close to the investment trigger. Therefore, predicting the aggregate response requires knowing how many firms are sitting near that boundary at the time of the policy.

### 9.44.9 Summary

The extension with persistent shocks yields five central lessons.

1. **Persistence matters for continuation values.** Because current shocks predict future shocks, the firm's decision becomes more forward-looking than in the i.i.d. benchmark.
2. **The normalized state $z = k/\varepsilon$ is the right object.** It transforms moving capital targets into fixed adjustment thresholds, producing a clean $(s, S)$-type policy.
3. **Depreciation creates endogenous drift.** Even without new shocks, non-adjusting firms move through the inaction region over time.

4. **Fixed costs generate overshooting.** After paying the fixed cost, the firm may optimally reset above the static frictionless target in order to postpone the next intervention.
5. **The distribution of firms becomes a macroeconomic state variable.** Aggregate responses to temporary investment subsidies depend on how many firms are close to the adjustment boundary.

**Final intuition.** At the firm level, the model is about when it is worth paying a fixed cost to realign capital with persistent profitability. At the aggregate level, the model is about how many firms are near that realignment margin. The first question gives us the policy function; the second gives us the business-cycle and policy implications.

# 10 Key Public Finance Results: Taxation, Incidence, and Monopoly

## 1. Deadweight loss from taxation

Consider a competitive market with inverse demand and inverse supply curves. A small specific tax $t$ creates a wedge between the consumer price $p_c$ and the producer price $p_p$, with

$$p_c - p_p = t.$$

The tax reduces the traded quantity from $Q^*$ to $Q_t$, and this quantity contraction generates a deadweight loss.

For a small tax, a standard local approximation gives

$$DWL \approx \frac{1}{2}\,\frac{\varepsilon^d \varepsilon^s}{\varepsilon^d + \varepsilon^s}\,\frac{Q^*}{P^*}\,t^2,$$

where $\varepsilon^d$ and $\varepsilon^s$ denote the absolute values of the demand and supply elasticities evaluated at the pre-tax equilibrium.

This formula delivers three central lessons:

1. Deadweight loss is approximately quadratic in the tax wedge.
2. Deadweight loss is larger when demand and supply are more elastic.
3. For a given revenue target, taxing relatively inelastic bases tends to generate smaller efficiency losses.

**Important qualification.** The previous intuition is fundamentally a *partial equilibrium* benchmark. In general equilibrium, the relevant object is not the elasticity of one isolated market, but the full system of compensated elasticities and cross-price effects. Taxing one good changes the demands for other goods, affects the labor–leisure margin, and interacts with the expenditure side of the government budget.

## 2. Statutory incidence versus economic incidence

A central distinction in public finance is between *statutory incidence* and *economic incidence.*

- **Statutory incidence** asks: who legally remits the tax?
- **Economic incidence** asks: who actually bears the burden once prices and quantities adjust?

In competitive markets, economic incidence is governed by elasticities: the less elastic side of the market bears more of the burden. Hence, who writes the check to the government is not the key economic question. What matters is how equilibrium prices adjust.

## 3. Why the inverse-elasticity intuition is only a benchmark

The statement

> "Tax goods with low elasticity."

is useful but incomplete.

In a general equilibrium Ramsey problem, the planner does not choose one tax in one market. The planner chooses an entire vector of taxes to raise revenue while minimizing the total welfare cost of distortions across all markets simultaneously. Three additional ideas matter:

1. **Cross-price effects.** Taxing good $i$ changes demand for goods $j \neq i$.
2. **Interaction with leisure.** If leisure cannot be taxed directly, commodity taxes can indirectly distort labor supply.
3. **Use of revenues.** If tax revenues finance public goods or government purchases, equilibrium effects depend on where that spending falls.

Thus, in general equilibrium, the relevant question is not whether the taxed market is inelastic in isolation, but whether taxing that base creates a small *total* distortion once all substitution margins are taken into account.

## 4. Corlett–Hague intuition

Suppose the government cannot tax leisure directly. Then a useful principle is:

> Tax relatively more the goods that are strong complements to leisure, and relatively less the goods that are strong complements to labor.

The intuition is that commodity taxation can be used to indirectly tax leisure. This is already a general equilibrium idea: what matters is not only own-price elasticity, but also how a good is related to the labor–leisure choice.

## 5. Monopoly and taxation

Now consider a monopolist facing inverse demand $P(Q)$, marginal revenue $MR(Q)$, and marginal cost $MC(Q)$.

Without taxes, the monopolist solves

$$MR(Q) = MC(Q).$$

**Specific (unit) tax.** If the government imposes a specific tax $t$ per unit sold, the monopolist solves

$$MR(Q) = MC(Q) + t.$$

Thus, a unit tax acts like an upward shift in marginal cost.

**Ad valorem tax.** If instead the government imposes an ad valorem tax rate $\tau$, so that the firm only keeps $(1-\tau)P(Q)$ per unit sold, the first-order condition becomes

$$(1-\tau)MR(Q) = MC(Q).$$

Equivalently, the monopolist behaves as if marginal revenue were rotated downward.

**Incidence under monopoly.** Under monopoly, pass-through is not governed by elasticities alone. It also depends on demand curvature and on the slope of marginal cost. Hence, unlike the competitive case, there is no simple universal formula saying that the less elastic side always bears the burden in the same way.

**A useful benchmark.** With linear demand and constant marginal cost, a specific tax is only partially passed through: the consumer price rises by $\frac{1}{2}t$. Hence monopoly does *not* generally pass the entire tax to consumers.

**Specific versus ad valorem taxation under monopoly.** For equal revenue, an ad valorem tax is often less distortionary than a specific tax under monopoly. The reason is subtle but important: a specific tax directly raises marginal cost unit by unit, while an ad valorem tax also taxes the monopoly markup. As a result, for a given revenue target, ad valorem taxation often implies a smaller quantity contraction.

## 6. Externalities and Pigouvian taxation

When a market generates an externality, the problem is no longer merely revenue extraction. There is also a corrective motive.

If marginal social cost exceeds marginal private cost by an amount $mec(Q)$, then the Pigouvian benchmark is

$$t^{Pigou}(Q) = mec(Q).$$

In that case, the tax is not mainly distorting the market away from efficiency. Instead, it corrects a pre-existing distortion by aligning private and social marginal cost.

For negative externalities, the logic is:

$$\text{private decision rule} \quad MC_p = MB$$

is replaced by

$$\text{socially efficient rule} \quad MC_p + MEC = MB.$$

## 7. Big picture

A good summary is the following:

- In **partial equilibrium**, low-elasticity tax bases are attractive because they generate smaller direct quantity distortions.
- In **general equilibrium**, one must consider the whole system of substitutions, labor supply, distributional objectives, and the use of public revenues.
- Under **market power**, tax incidence depends on pass-through, which is shaped by demand curvature and marginal cost.

- Under **externalities**, taxation may be corrective rather than merely distortionary.

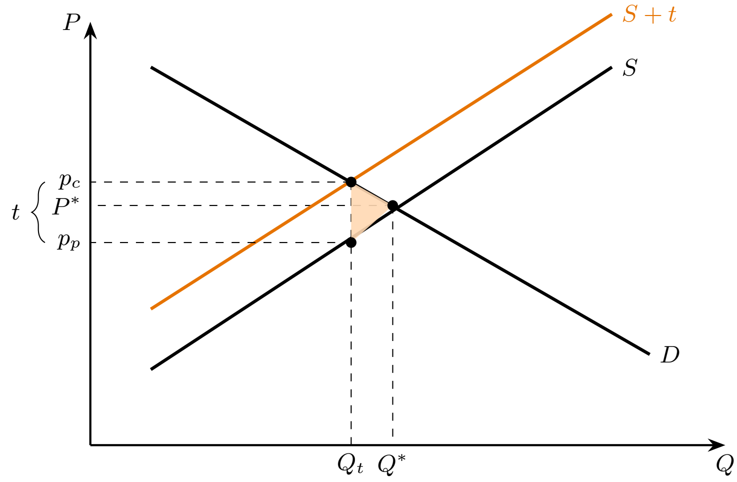

Figure 31: A specific tax in a competitive market. The shaded triangle is the deadweight loss.

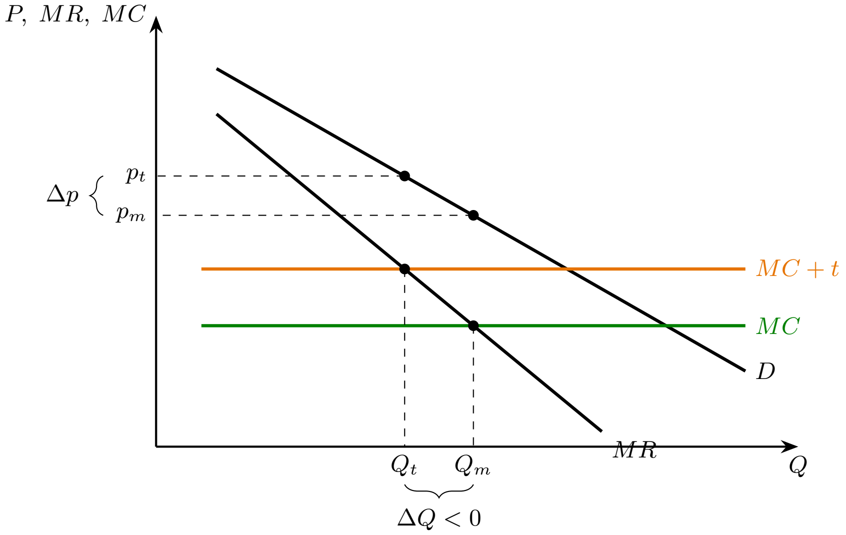

Figure 32: Specific taxation under monopoly: the tax shifts marginal cost upward, reduces output, and raises price.

**Remark.** The inverse-elasticity idea is best understood as a partial equilibrium benchmark. In general equilibrium, taxing one market does not only reduce consumption in that market. It also changes relative demands across all other goods, affects labor supply through the labor–leisure margin, and interacts with the way tax revenues are spent or rebated.

Hence, the relevant criterion in general equilibrium is not "which market is least elastic in isolation?" but rather "which tax system raises the required revenue with the smallest total distortion of the allocation?"

This is why Ramsey taxation is fundamentally a system-wide problem. The sufficient statistics are the full matrix of compensated demand elasticities, not a single own-price elasticity.

**Remark.** Under monopoly, tax incidence is shaped by pass-through. Pass-through depends not only on demand elasticity, but also on demand curvature and the slope of marginal cost. Thus, unlike the

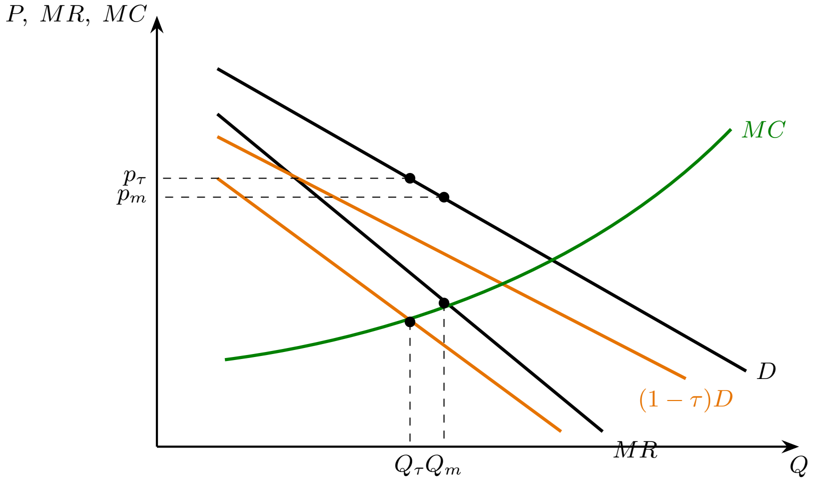

Figure 33: Ad valorem taxation under monopoly: the tax scales down effective demand and marginal revenue.

competitive benchmark, one cannot summarize incidence under monopoly with a simple "the less elastic side pays" rule.

# 11 Ricardian Equivalence

## Environment

Consider an endowment economy with a representative household, a government, and a one-period risk-free bond.

There is a single good each period. The household receives exogenous endowment $\{y_t\}_{t=0}^{\infty}$, consumes $\{c_t\}_{t=0}^{\infty}$, and trades one-period bonds $\{b_t\}_{t=0}^{\infty}$. The government purchases $\{g_t\}_{t=0}^{\infty}$, levies lump-sum taxes $\{T_t\}_{t=0}^{\infty}$, and issues bonds $\{B_t\}_{t=0}^{\infty}$.

Let the gross real interest rate between $t$ and $t+1$ be $1+r_t$. Define the gross discount factor

$$R_0 \equiv 1, \qquad R_t \equiv \prod_{s=0}^{t-1}(1+r_s) \quad \text{for } t \geq 1.$$

For simplicity, normalize initial private and public asset positions to zero:

$$b_{-1} = 0, \qquad B_{-1} = 0.$$

## Household problem

Given sequences $\{T_t\}_{t=0}^{\infty}$ and $\{r_t\}_{t=0}^{\infty}$, the household chooses $\{c_t, b_t\}_{t=0}^{\infty}$ to maximize

$$\sum_{t=0}^{\infty} \beta^t u(c_t)$$

subject to the flow budget constraints

$$c_t + b_t = y_t - T_t + (1+r_{t-1})b_{t-1}, \qquad t \geq 0.$$

Assuming the standard no-Ponzi condition, forward iteration yields the intertemporal budget constraint

$$c_0 + \sum_{t=1}^{\infty} \frac{c_t}{R_t} \leq (y_0 - T_0) + \sum_{t=1}^{\infty} \frac{y_t - T_t}{R_t}.$$

Equivalently,

$$\sum_{t=0}^{\infty} \frac{c_t}{R_t} \leq \sum_{t=0}^{\infty} \frac{y_t - T_t}{R_t}.$$

The Euler equation is

$$u'(c_t) = \beta(1 + r_t)u'(c_{t+1}).$$

## Government

The government's period-by-period budget constraint is

$$g_t + B_t = T_t + (1 + r_{t-1})B_{t-1}, \qquad t \geq 0.$$

Under the no-explosive-debt condition, forward iteration gives the government's present-value budget constraint

$$\sum_{t=0}^{\infty} \frac{T_t}{R_t} = \sum_{t=0}^{\infty} \frac{g_t}{R_t}.$$

Thus, once the path of government purchases $\{g_t\}$ is fixed, the government may choose the timing of taxes $\{T_t\}$, but not their present value.

## Market clearing

A competitive equilibrium consists of sequences

$$\{c_t, b_t, B_t, T_t, r_t\}_{t=0}^{\infty}$$

such that:

1. Given $\{T_t, r_t\}_{t=0}^{\infty}$, the household solves its optimization problem.
2. Given $\{r_t\}_{t=0}^{\infty}$ and exogenous $\{g_t\}_{t=0}^{\infty}$, government policy $\{T_t, B_t\}_{t=0}^{\infty}$ is feasible.
3. Markets clear for every $t$:

$$c_t + g_t = y_t, \qquad b_t + B_t = 0.$$

## Key observation

Substitute the government's present-value budget constraint into the household's intertemporal budget constraint:

$$\sum_{t=0}^{\infty} \frac{c_t}{R_t} \leq \sum_{t=0}^{\infty} \frac{y_t}{R_t} - \sum_{t=0}^{\infty} \frac{T_t}{R_t} = \sum_{t=0}^{\infty} \frac{y_t}{R_t} - \sum_{t=0}^{\infty} \frac{g_t}{R_t}.$$

Hence,

$$\sum_{t=0}^{\infty} \frac{c_t}{R_t} \leq \sum_{t=0}^{\infty} \frac{y_t - g_t}{R_t}.$$

This is the crucial step: the household's lifetime budget depends on the *present value* of taxes, not on their timing.

## Ricardian Equivalence proposition

**Proposition** (Ricardian Equivalence)**.** *Fix an exogenous path of government purchases* $\{g_t\}_{t=0}^{\infty}$*. Suppose*

$$\{\bar{c}_t, \bar{b}_t, \bar{B}_t, \bar{T}_t, \bar{r}_t\}_{t=0}^{\infty}$$

*is a competitive equilibrium.*

*Let* $\{\hat{T}_t\}_{t=0}^{\infty}$ *be any alternative tax sequence satisfying*

$$\sum_{t=0}^{\infty} \frac{\hat{T}_t}{R_t} = \sum_{t=0}^{\infty} \frac{\bar{T}_t}{R_t}.$$

*Then there exists an equilibrium*

$$\{\hat{c}_t, \hat{b}_t, \hat{B}_t, \hat{T}_t, \hat{r}_t\}_{t=0}^{\infty}$$

*such that*

$$\hat{c}_t = \bar{c}_t \qquad \text{and} \qquad \hat{r}_t = \bar{r}_t \quad \text{for all } t.$$

*That is, changing the timing of lump-sum taxes while keeping their present value fixed does not affect the equilibrium allocation.*

*Proof sketch.* Because

$$\sum_{t=0}^{\infty} \frac{\hat{T}_t}{R_t} = \sum_{t=0}^{\infty} \frac{\bar{T}_t}{R_t},$$

the household's intertemporal budget set is unchanged. Therefore, under the same interest-rate sequence, the same consumption plan remains optimal:

$$\hat{c}_t = \bar{c}_t \quad \forall t.$$

Now define household bond holdings recursively by

$$\hat{b}_t = y_t - \hat{T}_t + (1 + r_{t-1})\hat{b}_{t-1} - \hat{c}_t, \qquad \hat{b}_{-1} = 0.$$

Define government bond holdings recursively by

$$\hat{B}_t = \hat{T}_t - g_t + (1 + r_{t-1})\hat{B}_{t-1}, \qquad \hat{B}_{-1} = 0.$$

Using goods market clearing,

$$\hat{c}_t + g_t = y_t,$$

we obtain

$$\hat{b}_t + \hat{B}_t = y_t - \hat{c}_t - g_t - \hat{T}_t + \hat{T}_t + (1 + r_{t-1})(\hat{b}_{t-1} + \hat{B}_{t-1}).$$

Thus,

$$\hat{b}_t + \hat{B}_t = (1 + r_{t-1})(\hat{b}_{t-1} + \hat{B}_{t-1}).$$

Since $\hat{b}_{-1} + \hat{B}_{-1} = 0$, induction gives

$$\hat{b}_t + \hat{B}_t = 0 \quad \forall t.$$

Therefore, all markets clear and the new tax sequence supports the same allocation. □

## Economic intuition

Government debt is not net wealth for the representative household.

When the government cuts taxes today and finances that cut by issuing debt, the household receives a current transfer. But this transfer is exactly offset by the present value of higher future taxes required to service and repay the debt. Hence the household's lifetime resources are unchanged.

In equilibrium, the household simply saves the tax cut in order to pay those future taxes. Aggregate consumption does not change.

A useful way to say it is:

$$\text{public debt} = \text{private asset} - \text{future tax liability}.$$

Once future tax liabilities are fully internalized, debt does not create additional net wealth.

## Why the assumptions matter

**Remark** (Lump-sum taxes)**.** This assumption is doing essential work. Lump-sum taxes do not distort labor supply, saving, or any other margin. Therefore, changing the timing of taxation changes only the intertemporal distribution of payments, not incentives.

With distortionary taxes, timing matters. For example, replacing current labor taxes with future labor taxes changes labor wedges across dates, so the allocation generally changes.

**Remark** (No borrowing constraints)**.** If households are liquidity constrained, a tax cut today can raise current consumption even if future taxes increase by exactly the same present value. In that case, the household cannot perfectly smooth consumption by borrowing against future income, so the timing of taxes matters.

**Remark** (Infinite lives versus bequests)**.** The assumption that households are infinitely lived guarantees that they internalize the entire future tax path.

A more general route to the same conclusion is a dynastic model with finitely lived agents connected by operative altruistic bequests. In that case, parents care about their descendants and leave transfers that offset future taxes.

Without infinite lives or operative bequests, government debt can redistribute across generations, and Ricardian equivalence typically fails.

**Remark** (Representative-agent structure)**.** The one-household setup shuts down redistribution across agents. With heterogeneity, debt-financed tax cuts may transfer resources toward households with high marginal propensities to consume, so aggregate consumption may respond even if the present value of taxes is unchanged.

## Bottom line

Ricardian equivalence is not the claim that "debt never matters." Rather, it is the claim that in a very special benchmark economy, the *timing* of lump-sum taxes is irrelevant for allocations once the present value of government purchases is fixed.

What matters for the household is not whether taxes are collected today or tomorrow, but the present value of taxes it must ultimately pay.

# 12 Tax Smoothing (Barro, 1979)

The core idea of Barro's tax-smoothing model is that the government chooses the timing of taxes so as to minimize the present discounted value of the distortionary costs of taxation, subject to its intertemporal budget constraint.

The logic parallels the permanent-income hypothesis. A household smooths consumption because marginal utility is diminishing. Here, the government smooths tax rates because the marginal distortion from taxation is increasing.

**Distortionary cost of taxation.** Let $T_t$ denote total tax revenues at date $t$, and let $Y_t$ denote output. Define the average tax rate by

$$\tau_t \equiv \frac{T_t}{Y_t}.$$

Assume that the per-period distortionary cost of taxation is

$$Y_t\, \phi\!\left(\frac{T_t}{Y_t}\right) = Y_t\, \phi(\tau_t),$$

where $\phi(\cdot)$ satisfies

$$\phi(0) = 0, \qquad \phi'(0) = 0, \qquad \phi''(\tau) > 0.$$

Thus, distortionary costs are convex in the tax rate: raising a given amount of revenue is increasingly costly at higher tax rates.

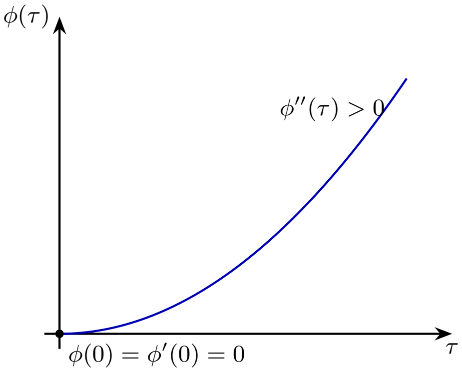

Figure 34: Convex distortionary cost of taxation.

**Government problem.** Take as given an exogenous sequence of government expenditures $\{g_t\}_{t\geq 0}$. The government chooses taxes $\{T_t\}_{t\geq 0}$ and one-period debt $\{B_t\}_{t\geq 0}$ to solve

$$\min_{\{T_{t+j}, B_{t+j}\}_{j\geq 0}} \; E_t \sum_{j=0}^{\infty} \beta^j Y_{t+j} \phi\!\left(\frac{T_{t+j}}{Y_{t+j}}\right)$$

subject to the flow budget constraints

$$g_{t+j} + (1 + r_{t+j-1})B_{t+j-1} = T_{t+j} + B_{t+j}, \qquad j \geq 0.$$

**Lagrangian.** It is convenient to write the problem with multipliers $\{\lambda_{t+j}\}_{j\geq 0}$:

$$\mathcal{L} = E_t \sum_{j=0}^{\infty} \beta^j \left[ Y_{t+j} \phi\!\left(\frac{T_{t+j}}{Y_{t+j}}\right) + \lambda_{t+j}\Big(g_{t+j} + (1 + r_{t+j-1})B_{t+j-1} - T_{t+j} - B_{t+j}\Big)\right].$$

**First-order conditions.** Differentiating with respect to $T_{t+j}$ gives

$$\phi'(\tau_{t+j}) = \lambda_{t+j},$$

because

$$\frac{\partial}{\partial T_{t+j}}\left[Y_{t+j}\phi\left(\frac{T_{t+j}}{Y_{t+j}}\right)\right] = \phi'(\tau_{t+j}).$$

Differentiating with respect to $B_{t+j}$ gives

$$\lambda_{t+j} = \beta(1 + r_{t+j})E_t[\lambda_{t+j+1}].$$

Combining the two conditions yields the Euler equation for optimal taxation:

$$\phi'(\tau_t) = \beta(1 + r_t)E_t[\phi'(\tau_{t+1})]\,.$$

**Remark.** This is the tax-smoothing condition. The government chooses the time profile of taxes so that the current marginal distortionary cost of taxation equals the discounted expected marginal distortionary cost tomorrow.

**Special case: quadratic costs.** Suppose

$$\phi(\tau) = \frac{\bar{\phi}}{2}\tau^2, \qquad \bar{\phi} > 0.$$

Then

$$\phi'(\tau) = \bar{\phi}\tau,$$

so the Euler equation becomes

$$\bar{\phi}\tau_t = \beta(1 + r_t)E_t[\bar{\phi}\tau_{t+1}],$$

or equivalently,

$$\tau_t = \beta(1 + r_t)E_t[\tau_{t+1}].$$

If, in addition,

$$\beta(1 + r_t) = 1,$$

then

$$\tau_t = E_t[\tau_{t+1}].$$

**Remark.** Under certainty, the condition above implies a constant tax rate over time:

$$\tau_t = \tau_{t+1} = \tau_{t+2} = \cdots$$

Hence the government uses debt to absorb temporary fluctuations in spending, rather than moving tax rates one-for-one with expenditures.

**Economic intuition.** Because distortionary costs are convex, it is inefficient to finance a temporary increase in government spending by sharply increasing taxes only in the current period. A better policy is to spread the tax burden over time through debt issuance.

This implies:

1. temporary increases in $g_t$ should mainly be financed with deficits;

2. persistent or permanent increases in $g_t$ should lead to higher tax rates;
3. in the benchmark case, tax rates are smoother than government spending.

**Connection with the permanent-income logic.** The analogy with consumption smoothing is exact:

- households smooth consumption because utility is concave;
- governments smooth taxes because distortionary costs are convex.

So the tax-smoothing model can be read as a public-finance analogue of the permanent-income hypothesis.

## 12.1 Aiyagari, Marcet, Sargent, and Seppälä (2002): tax smoothing without state-contingent debt

In the benchmark tax-smoothing problem, and under the special case

$$\beta(1+r_t) = 1,$$

the Euler equation for optimal taxation becomes

$$\phi'(\tau_t) = E_t[\phi'(\tau_{t+1})] \, .$$

Thus, the *marginal distortionary cost* of taxation is a martingale.

**A Jensen-based interpretation.** Suppose that $\phi''(\tau) > 0$ and $\phi'''(\tau) > 0$, so that $\phi'(\tau)$ is increasing and strictly convex. Then Jensen's inequality implies

$$E_t[\phi'(\tau_{t+1})] \geq \phi'(E_t[\tau_{t+1}]) \, ,$$

with strict inequality whenever $\tau_{t+1}$ is nondegenerate. Using the Euler equation,

$$\phi'(\tau_t) = E_t[\phi'(\tau_{t+1})] \geq \phi'(E_t[\tau_{t+1}]) \, .$$

Since $\phi'$ is increasing, it follows that

$$\tau_t \geq E_t[\tau_{t+1}],$$

with strict inequality if uncertainty is nontrivial and $\phi'$ is strictly convex. Hence the tax rate itself is a *supermartingale.*

**Remark.** The exact martingale object is $\phi'(\tau_t)$, not $\tau_t$. Under uncertainty and convex marginal deadweight costs, the tax rate drifts downward in expectation even though the marginal distortion is smoothed.

**Local approximation.** A second-order Taylor expansion of $\phi'(\tau_{t+1})$ around $\tau_t$ gives

$$\phi'(\tau_{t+1}) \approx \phi'(\tau_t) + \phi''(\tau_t)(\tau_{t+1} - \tau_t) + \frac{1}{2}\phi'''(\tau_t)(\tau_{t+1} - \tau_t)^2.$$

Taking conditional expectations and using

$$\phi'(\tau_t) = E_t[\phi'(\tau_{t+1})],$$

we obtain

$$0 \approx \phi''(\tau_t)\big(E_t[\tau_{t+1}] - \tau_t\big) + \frac{1}{2}\phi'''(\tau_t)E_t\left[(\tau_{t+1} - \tau_t)^2\right].$$

Therefore,

$$E_t[\tau_{t+1}] \approx \tau_t - \frac{1}{2}\frac{\phi'''(\tau_t)}{\phi''(\tau_t)}E_t\left[(\tau_{t+1} - \tau_t)^2\right].$$

If $\phi''(\tau_t) > 0$ and $\phi'''(\tau_t) > 0$, then

$$E_t[\tau_{t+1}] < \tau_t.$$

**Remark.** This approximation makes the precautionary logic transparent. For a given mean tax rate, uncertainty raises expected marginal deadweight losses because $\phi'$ is convex. To offset that effect and keep the Euler equation valid, the planner lowers the conditional mean of future tax rates.

**Economic intuition.** When the government cannot trade state-contingent debt, it cannot perfectly insure fiscal shocks. Instead, it uses risk-free assets and liabilities to self-insure. This creates a precautionary motive for the government to accumulate assets in good times, so as to reduce the need for high distortionary taxation in bad times. As a result, optimal tax rates tend to drift downward over time.

The key message is not simply that "taxes go to zero" mechanically from Jensen's inequality alone. Rather, the combination of incomplete fiscal insurance and convex marginal deadweight losses induces a policy of gradual asset accumulation that pushes taxes downward in expectation.

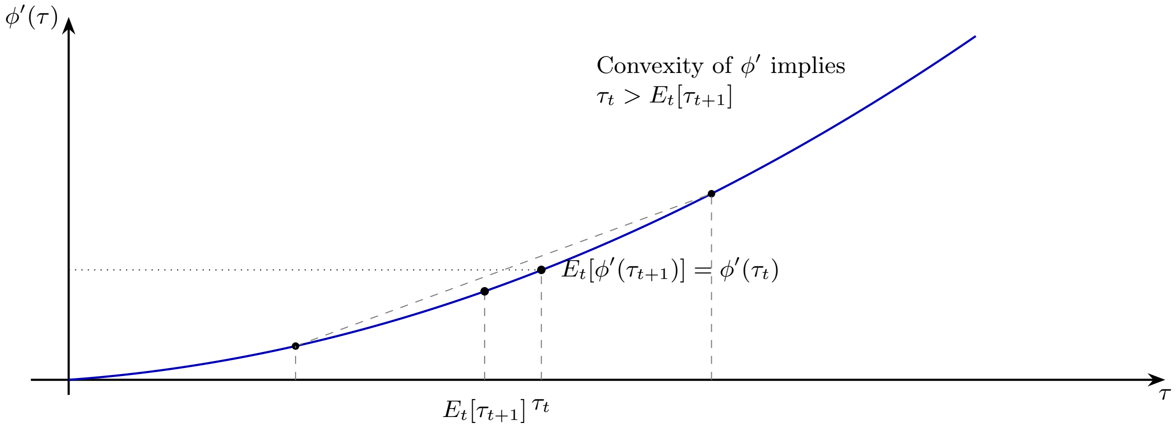

Figure 35: With $\phi'(\tau)$ convex, smoothing marginal distortions implies downward drift in the tax rate itself.

**A variance-based intuition: mean-preserving spreads.** There is a complementary way to understand the AMSS result. Suppose that the government keeps the conditional mean of next period's tax rate fixed,

$$E_t[\tau_{t+1}] = \bar{\tau},$$

but faces a higher conditional variance of $\tau_{t+1}$. That is, the distribution of $\tau_{t+1}$ becomes a mean-preserving spread.

Because $\phi'(\tau)$ is increasing and convex when

$$\phi''(\tau) > 0 \qquad \text{and} \qquad \phi'''(\tau) > 0,$$

Jensen's inequality implies

$$E_t[\phi'(\tau_{t+1})] > \phi'(E_t[\tau_{t+1}]).$$

Moreover, for a fixed mean $E_t[\tau_{t+1}]$, a larger variance raises

$$E_t[\phi'(\tau_{t+1})].$$

Using the Euler equation

$$\phi'(\tau_t) = E_t[\phi'(\tau_{t+1})],$$

we conclude that an increase in the conditional variance of future tax rates must raise the current marginal distortionary cost. Since $\phi'$ is increasing, this implies a higher current tax rate:

$$\mathrm{Var}_t(\tau_{t+1}) \uparrow \implies E_t[\phi'(\tau_{t+1})] \uparrow \implies \phi'(\tau_t) \uparrow \implies \tau_t \uparrow .$$

**Remark.** Holding fixed the conditional mean of future tax rates, more fiscal risk makes future marginal deadweight losses more expensive in expectation. The government responds by taxing more today, or equivalently by accumulating more assets today, in order to reduce its exposure to high-tax states tomorrow.

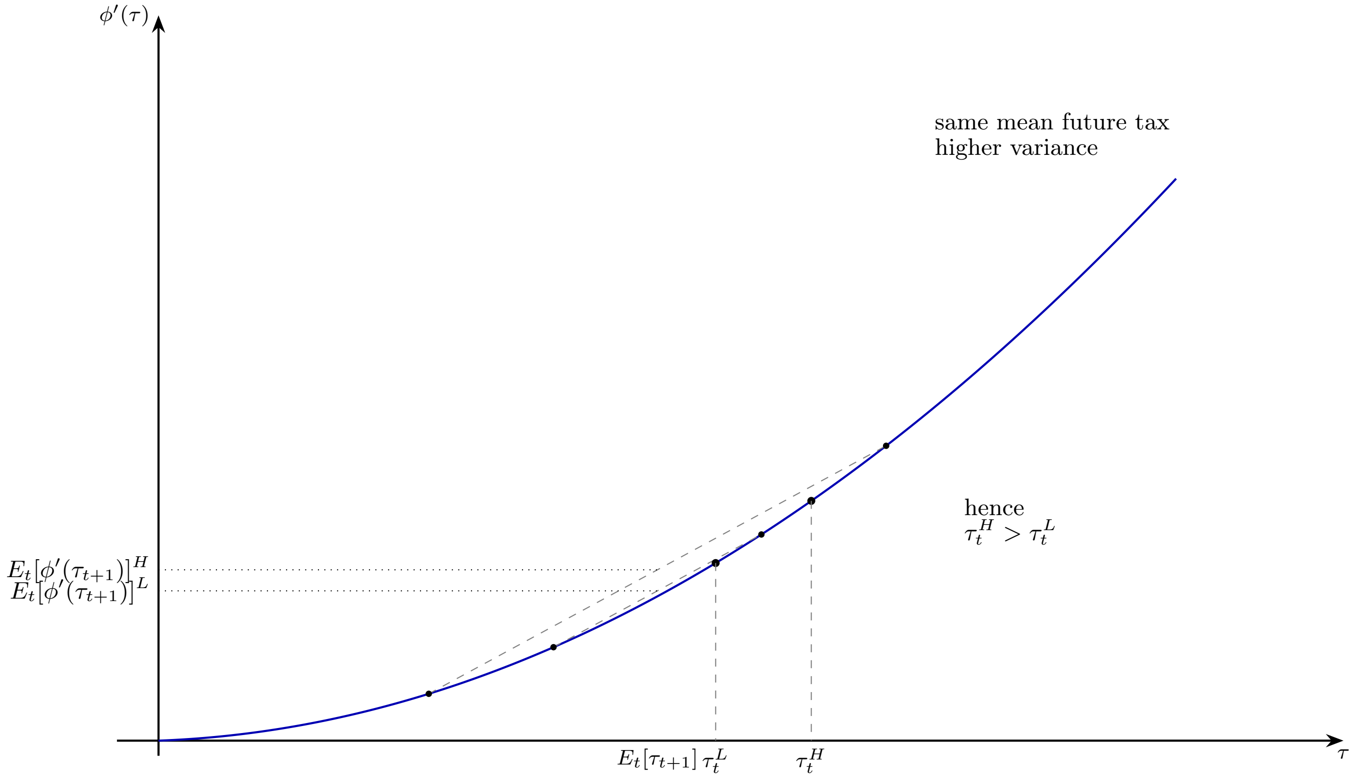

Figure 36: Holding fixed $E_t[\tau_{t+1}]$, a mean-preserving spread raises $E_t[\phi'(\tau_{t+1})]$. By the Euler equation, this increases the current optimal tax rate.

## 12.2 Capital taxation with heterogeneous agents

We now move to a heterogeneous-agent environment in which capital taxation acquires a redistributive role.

**Households.** There are two types of households, $j \in \{A, B\}$, with lifetime utility

$$\sum_{t=0}^{\infty} \beta^t u^j(c_t^j, n_t^j), \qquad j \in \{A, B\}.$$

Type $A$ households are hand-to-mouth workers: they do not save and do not accumulate capital. Their budget constraint is

$$c_t^A = (1 - \tau_t^n) w_t n_t^A + T_t^A.$$

Type $B$ households receive labor income and can save in both physical capital and one-period bonds. Their budget constraint is

$$c_t^B + k_{t+1}^B + b_t^B = (1 - \tau_t^k) R_t k_t^B + (1 - \delta) k_t^B + (1 - \tau_t^n) w_t n_t^B + T_t^B + (1 + r_{t-1}) b_{t-1}^B.$$

**Household optimality conditions.** Let

$$u_{c,t}^j \equiv \frac{\partial u^j(c_t^j, n_t^j)}{\partial c}, \qquad u_{n,t}^j \equiv \frac{\partial u^j(c_t^j, n_t^j)}{\partial n}.$$

For both types, the labor supply condition is

$$u_{n,t}^j + (1 - \tau_t^n) w_t u_{c,t}^j = 0, \qquad j \in \{A, B\}.$$

For type $B$, optimal bond accumulation implies

$$u_{c,t}^B = \beta (1 + r_t) E_t \left[ u_{c,t+1}^B \right].$$

Optimal capital accumulation implies

$$u_{c,t}^B = \beta E_t \left[ u_{c,t+1}^B \left( (1 - \tau_{t+1}^k) R_{t+1} + 1 - \delta \right) \right].$$

Combining the two Euler equations yields the no-arbitrage condition

$$1 + r_t = (1 - \tau_{t+1}^k) R_{t+1} + 1 - \delta,$$

or equivalently,

$$(1 - \tau_{t+1}^k) R_{t+1} = r_t + \delta.$$

**Remark.** This expression is very informative. A capital income tax drives a wedge between the after-tax return on physical capital and the return on the risk-free bond, net of depreciation.

**Firms.** Competitive firms solve

$$\max_{K_t, N_t} \; F(K_t, N_t) - w_t N_t - R_t K_t.$$

Hence factor prices satisfy

$$F_N(K_t, N_t) = w_t, \qquad F_K(K_t, N_t) = R_t.$$

**Government budget constraint.** Taking government purchases $\{g_t\}_{t\geq 0}$ as given, the government chooses labor taxes, capital taxes, transfers, and debt subject to

$$g_t + B_t = \tau_t^k R_t K_t + \tau_t^n w_t N_t + (1 + r_{t-1}) B_{t-1} - T_t^A - T_t^B.$$

**Market clearing.** Aggregate feasibility requires

$$c_t^A + c_t^B + i_t + g_t = y_t = F(K_t, N_t),$$

with capital accumulation

$$i_t = K_{t+1} - (1 - \delta) K_t.$$

**Government objective.** A natural Ramsey objective is

$$\max_{\{\tau_t^n, \tau_t^k, T_t^A, T_t^B, B_t\}_{t\geq 0}} \sum_{t=0}^{\infty} \beta^t \left[ \alpha u^A(c_t^A, n_t^A) + (1 - \alpha) u^B(c_t^B, n_t^B) \right],$$

subject to household optimality, firm optimality, market clearing, and the government budget constraint.

**Main intuition.** In a representative-agent benchmark, a capital tax is purely distortionary along the intertemporal margin, which is why the long-run optimal capital tax often tends to zero. Here, however, capital taxation also has a redistributive role.

Type $A$ households live hand-to-mouth and do not own assets. Type $B$ households own capital and bonds. Therefore, a capital income tax becomes a way of taxing the asset-rich households and transferring resources toward households that cannot self-insure through savings.

This changes the Ramsey logic fundamentally: capital taxation is no longer only about intertemporal efficiency; it is also about redistribution and insurance across heterogeneous agents.

**Remark.** The planner now faces a genuine tradeoff: a higher capital tax worsens intertemporal distortions for savers, but it may improve welfare by redistributing toward constrained households. That tradeoff is precisely what is absent in the representative-agent zero-capital-tax benchmark.

**A primal reformulation with after-tax prices.** To make the Ramsey problem more transparent, it is convenient to work with the after-tax rental rate on capital and the after-tax wage:

$$\hat{R}_t \equiv (1 - \tau_t^k) R_t, \qquad \hat{w}_t \equiv (1 - \tau_t^n) w_t.$$

Since only type $B$ households save, aggregate capital is owned by type $B$:

$$K_t = k_t^B, \qquad N_t = n_t^A + n_t^B.$$

Using these definitions, government revenues can be written as

$$\text{Rev}_t = \tau_t^k R_t K_t + \tau_t^n w_t N_t = K_t (R_t - \hat{R}_t) + N_t (w_t - \hat{w}_t).$$

Under perfect competition, firms satisfy

$$R_t = F_K(K_t, N_t), \qquad w_t = F_N(K_t, N_t).$$

If $F$ is constant returns to scale, Euler's theorem implies

$$F(K_t, N_t) = F_K(K_t, N_t)K_t + F_N(K_t, N_t)N_t = R_t K_t + w_t N_t,$$

so revenues admit the useful alternative expression

$$\mathrm{Rev}_t = F(K_t, N_t) - \hat{R}_t K_t - \hat{w}_t N_t.$$

This is the key step behind the primal representation: taxes enter only through the wedges $\hat{R}_t$ and $\hat{w}_t$, rather than through $\tau_t^k$ and $\tau_t^n$ separately.

**Remark.** The labor wedge applies at the micro level to both household types. Hence the aggregate term $\hat{w}_t N_t$ is not merely shorthand: it is exactly the sum of after-tax labor income across types,

$$\hat{w}_t N_t = \hat{w}_t n_t^A + \hat{w}_t n_t^B.$$

**Planner's problem and Lagrangian.** A convenient Ramsey problem is then

$$\max_{\{c_t^A, c_t^B, n_t^A, n_t^B, K_{t+1}, B_t, \hat{w}_t, \hat{R}_t, T_t^A, T_t^B\}_{t\geq 0}} \sum_{t=0}^{\infty} \beta^t \Big[\alpha u^A(c_t^A, n_t^A) + (1-\alpha) u^B(c_t^B, n_t^B)\Big],$$

subject to the equilibrium conditions written in terms of allocations and wedges. One useful Lagrangian is[6]

$$\begin{aligned}
\mathcal{L} = \sum_{t=0}^{\infty} \beta^t \Bigg\{ & \alpha u^A(c_t^A, n_t^A) + (1-\alpha) u^B(c_t^B, n_t^B) \\
& + \sum_{j \in \{A,B\}} \theta_t^j \Big(u_{n,t}^j + \hat{w}_t u_{c,t}^j\Big) \\
& + \gamma_t \Big[F(K_t, N_t) - \hat{R}_t K_t - \hat{w}_t N_t + (\hat{R}_t + 1 - \delta) B_{t-1} - T_t^A - T_t^B - B_t - g_t\Big] \\
& + \mu_t \Big[F(K_t, N_t) - c_t^A - c_t^B - K_{t+1} + (1-\delta) K_t - g_t\Big] \\
& + \xi_t \Big[\hat{w}_t n_t^A + T_t^A - c_t^A\Big] + \varphi_t \Big[b_t^B + B_t\Big] \\
& + \eta_t \Big[u_{c,t}^B - \beta u_{c,t+1}^B \big(\hat{R}_{t+1} + 1 - \delta\big)\Big] \Bigg\}.
\end{aligned}$$

Here:

- $\theta_t^A$ and $\theta_t^B$ impose the two labor-supply conditions,

$$u_{n,t}^j + \hat{w}_t u_{c,t}^j = 0, \qquad j \in \{A, B\};$$

- $\gamma_t$ imposes the government budget constraint in primal form;
- $\mu_t$ imposes aggregate feasibility;
- $\xi_t$ imposes the hand-to-mouth budget constraint of type $A$,

$$c_t^A = \hat{w}_t n_t^A + T_t^A;$$

---

[6] Remember that $N_t = n_t^A + n_t^B$.

- $\varphi_t$ imposes bond-market clearing,
$$b_t^B + B_t = 0;$$
- $\eta_t$ imposes the Euler equation of the saving household,
$$u_{c,t}^B = \beta u_{c,t+1}^B \big(\hat{R}_{t+1} + 1 - \delta\big).$$

**A Chamley–Judd-type steady-state argument.** The key first-order condition for capital accumulation is the one with respect to $K_{t+1}$:

$$-\mu_t + \beta\Big[\gamma_{t+1}\big(F_{K,t+1} - \hat{R}_{t+1}\big) + \mu_{t+1}\big(F_{K,t+1} + 1 - \delta\big)\Big] = 0.$$

Equivalently,
$$\mu_t = \beta\Big[\gamma_{t+1}\big(F_{K,t+1} - \hat{R}_{t+1}\big) + \mu_{t+1}\big(F_{K,t+1} + 1 - \delta\big)\Big].$$

Suppose the economy converges to an interior steady state in which the multipliers also converge:

$$\mu_t \to \mu, \qquad \gamma_t \to \gamma, \qquad F_{K,t} \to F_K, \qquad \hat{R}_t \to \hat{R}.$$

Then the previous condition becomes

$$\mu = \beta\Big[\gamma(F_K - \hat{R}) + \mu(F_K + 1 - \delta)\Big].$$

Rearranging,
$$\mu = \beta(\gamma + \mu)F_K - \beta\gamma\hat{R} + \beta\mu(1 - \delta).$$

Now use the steady-state Euler equation of the saving household:

$$1 = \beta\big(\hat{R} + 1 - \delta\big).$$

Add and subtract $\beta\mu\hat{R}$ in the previous expression:

$$\mu = \beta(\gamma + \mu)F_K - \beta(\gamma + \mu)\hat{R} + \beta\mu\big(\hat{R} + 1 - \delta\big).$$

Using $1 = \beta(\hat{R} + 1 - \delta)$, we obtain

$$\mu = \beta(\gamma + \mu)(F_K - \hat{R}) + \mu.$$

Hence
$$\beta(\gamma + \mu)(F_K - \hat{R}) = 0.$$

Provided $\gamma + \mu > 0$, it follows that
$$F_K = \hat{R}.$$

But $\hat{R} = (1 - \tau^k)R$ and, in competitive equilibrium, $R = F_K$. Therefore

$$(1 - \tau^k)F_K = F_K \qquad \Longrightarrow \qquad \tau^k = 0.$$

**Remark.** This is the same basic long-run logic as in the Chamley–Judd result: a permanent tax on reproducible capital distorts the intertemporal margin forever, so in an interior steady state the Ramsey planner eliminates that wedge. Heterogeneity changes the transition dynamics and creates a redistributive motive, but it does not by itself justify a permanent steady-state tax on reproducible capital.

**Why time 0 is special.** The previous argument is about the limiting tax on *reproducible* capital. It does *not* imply that the tax on capital owned at date 0 must be zero. At $t = 0$, the stock $K_0$ is predetermined, so its supply is perfectly inelastic. Therefore a tax on installed capital at date 0 is nondistortionary at the margin: it does not affect the quantity of preexisting capital.

This is why a Ramsey planner may want to impose a large one-time capital levy at the initial date. Economically, this is closer to a tax on wealth already installed than to a tax on the future return to new capital accumulation.

**Remark.** The relevant supply elasticity at $t = 0$ is *zero*, not infinity. A vertical supply curve means quantity does not respond to the tax-inclusive return. By contrast, an infinite elasticity would correspond to a perfectly elastic supply curve.

**A note on the first-order condition for $\hat{R}_t$ and public debt.** Differentiating the Lagrangian with respect to $\hat{R}_{t+1}$ yields a condition of the form

$$\gamma_t(B_t - K_{t+1}) = \beta\eta_t u^B_{c,t+1}.$$

In a steady state this becomes

$$\gamma(B - K) = \beta\eta u^B_c.$$

This should not be read as saying that the government "prefers bonds" to capital. Rather, it links three shadow objects: the multiplier on the government budget constraint ($\gamma$), the multiplier on the implementability/Euler restriction ($\eta$), and the gap between public debt and private capital holdings. So the condition is a consistency relation among wedges and shadow values, not a standalone debt-demand equation.

**Figures.** The two graphs below summarize the two key ideas: a large initial capital levy is possible because installed capital is predetermined, while the long-run tax on reproducible capital vanishes.

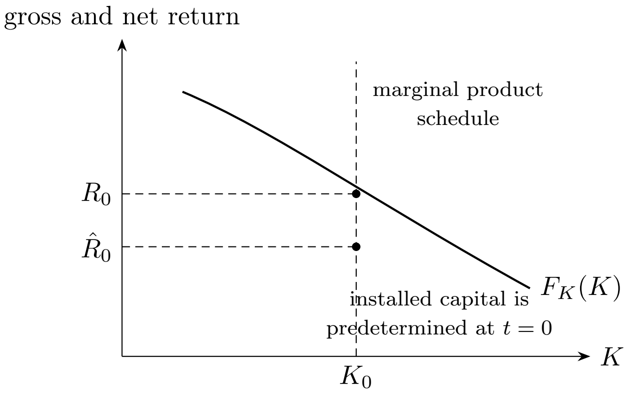

Figure 37: At the initial date, the stock of installed capital is predetermined. A one-time capital levy lowers the net return received by asset holders from $R_0$ to $\hat{R}_0$, but it does not change the quantity $K_0$.

**Bottom line.** With heterogeneous agents, capital taxation has a genuine redistributive role along the transition. But if the economy converges to an interior steady state and the planner can use other intertemporal instruments, the long-run capital tax on reproducible capital still converges to zero, while the initial tax on preexisting capital can be positive and large.

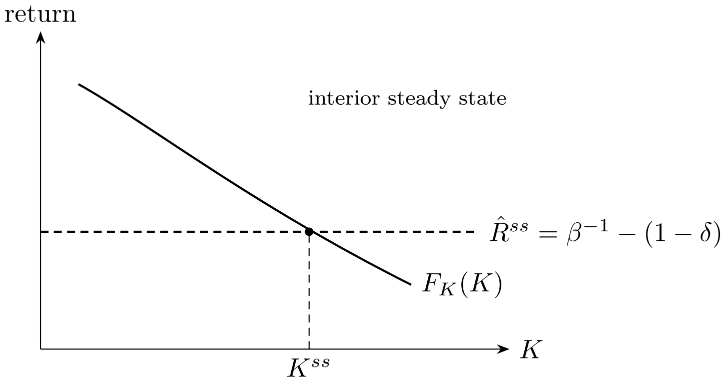

Figure 38: In the interior steady state, the Euler equation implies $\hat{R}^{ss} = \beta^{-1} - (1-\delta)$. The Ramsey first-order condition then implies $F_K(K^{ss}) = \hat{R}^{ss}$, so the steady-state capital tax wedge is zero.

**Remark** (An important detail)**.** A useful way to close these Ramsey-taxation results is to distinguish sharply between the *complete-markets* benchmark and the case in which the government can trade only *risk-free* assets.

With state-contingent debt, the government can insure directly across states: high-expenditure realizations can be financed by state-contingent payoffs rather than by large contemporaneous tax changes. For that reason, the current value of government debt depends only on the current state and expected future expenditures, not on the detailed past history. In that environment, there is no strong need for tax distortions to inherit the effects of old shocks.

By contrast, when the government can issue only risk-free debt, it loses that insurance margin. It can smooth taxes over time, but it cannot fully smooth them across states. Intuitively, the government is then in a position analogous to an agent with incomplete insurance who can save only through a non-state- contingent asset. As a result, shocks to government expenditures leave persistent effects on future debt and tax rates. This is the source of the martingale/supermartingale logic emphasized in the Aiyagari–Marcet–Sargent–Seppälä environment.

This helps unify the results of the last two classes. In the AMSS case, incomplete insurance across states generates history dependence and long-run effects of fiscal shocks on taxes and debt. In the Chamley–Judd logic, by contrast, the key force is that a permanent tax on reproducible capital imposes a permanent intertemporal distortion, so in an interior steady state the optimal capital tax vanishes asymptotically. Thus, both results reflect tax-smoothing logic, but they operate on different margins: AMSS is about persistence created by missing state-contingent assets, whereas Chamley–Judd is about the long-run elimination of the capital wedge.

# 13 Overlapping Generations Model

Following Samuelson (1958) and Diamond (1965), generations live for two periods: young and old. At date $t$, a young household values consumption when young and when old according to

$$u(c_t^y, c_{t+1}^o),$$

and is endowed with

$$(\omega_t^y, \omega_{t+1}^o).$$

A young household at date $t$ solves

$$\max_{b_t} \; u(c_t^y, c_{t+1}^o)$$

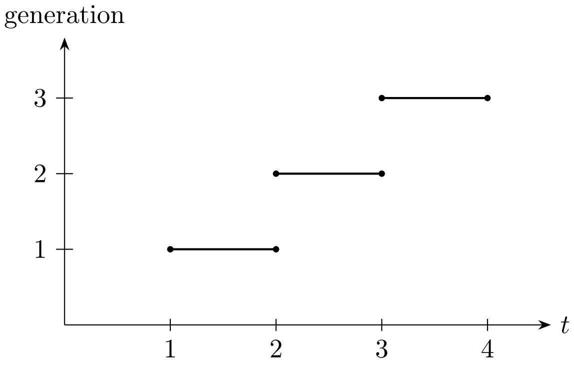

Figure 39: Each cohort lives for two periods.

subject to

$$c_t^y + b_t = \omega_t^y - T_t^y, \qquad c_{t+1}^o = \omega_{t+1}^o - T_{t+1}^o + (1 + r_t)b_t.$$

The initial old are special: they enter the economy already old and only consume at date 0.

We assume throughout that consumers maximize, markets clear, the government satisfies its budget constraint, and cohort sizes are constant:

$$N_t = 1.$$

**Special case.** Consider

$$u(c_t^y, c_{t+1}^o) = \log c_t^y + \beta \log c_{t+1}^o, \qquad (\omega_t^y, \omega_{t+1}^o) = (\bar{\omega} - \varepsilon, \varepsilon).$$

**Autarky.** Under autarky there is no intergenerational trade, so the young consume their young endowment and the old consume their old endowment:

$$c_t^y = \bar{\omega} - \varepsilon, \qquad c_{t+1}^o = \varepsilon.$$

The Euler equation is

$$\frac{u_{c_t^y}}{u_{c_{t+1}^o}} = 1 + r_t.$$

Under the log specification this becomes

$$\frac{c_{t+1}^o}{\beta c_t^y} = 1 + r_t,$$

or equivalently

$$\frac{c_{t+1}^o}{c_t^y} = \beta(1 + r_t).$$

Evaluated at the autarkic allocation,

$$1 + r_t = \frac{\varepsilon}{\beta(\bar{\omega} - \varepsilon)}.$$

Hence the autarkic gross interest rate is pinned down by the steepness of the endowment profile over the life cycle: when endowments are very backloaded toward youth, the equilibrium return is low.

**Why autarky may fail to be Pareto efficient.** If

$$1 + r_t < 1,$$

the competitive allocation is not Pareto efficient. The key intuition is that one unit transferred to the old today requires taking back less than one unit from the next young. This opens the door to a Pareto improvement: the current old are made strictly better off, and future young can be compensated at a lower resource cost than the initial transfer.

For intuition, fix total lifetime resources at $\bar{\omega}$. When $\beta = 1$, utility along the feasibility set

$$c_t^y + c_{t+1}^o = \bar{\omega}$$

is

$$W(c_t^y) = \log c_t^y + \log(\bar{\omega} - c_t^y),$$

which is strictly concave and reaches its maximum at

$$c_t^y = c_{t+1}^o = \frac{\bar{\omega}}{2}.$$

Thus, if the endowment point is $(\bar{\omega} - \varepsilon, \varepsilon)$ with $\varepsilon < \bar{\omega}/2$, the allocation is too uneven across the life cycle.

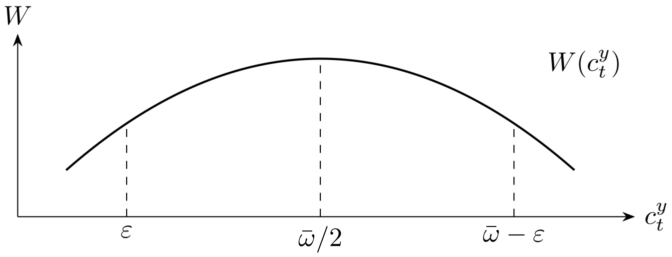

Figure 40: For $\beta = 1$, utility along the feasible set is maximized by smoothing consumption across the two dates.

More generally, when $\beta \neq 1$, the efficient allocation does not equalize consumption one-for-one across dates. Instead, it satisfies

$$\frac{1}{c_t^y} = \beta \frac{1}{c_{t+1}^o} \qquad \Longleftrightarrow \qquad c_{t+1}^o = \beta c_t^y.$$

So the equal split $\bar{\omega}/2$ is best viewed as the $\beta = 1$ benchmark that makes the inefficiency especially transparent.

**Government borrowing.** Now specialize further to the case

$$\beta = 1,$$

and suppose the government issues one-period debt, with

$$B_1 < 0,$$

so that the government is a net borrower in the first period.

For simplicity, assume that from period $t \geq 1$ onward there are no taxes on either the young or the old:

$$T_t^y = T_t^o = 0 \qquad \text{for all } t \geq 1.$$

A generation born at date $t$ then faces lifetime resources

$$c_t^y + \frac{c_{t+1}^o}{1 + r_t} = \bar{\omega} - \varepsilon + \frac{\varepsilon}{1 + r_t}.$$

Since $\beta = 1$, the household wants to smooth consumption across the two dates as much as prices permit, and the Euler equation implies

$$c_t^y = \frac{c_{t+1}^o}{1+r_t}.$$

Combining this with the lifetime budget constraint gives

$$c_t^y = \frac{1}{2}\left(\bar{\omega} - \varepsilon + \frac{\varepsilon}{1+r_t}\right).$$

Hence the savings of the young are

$$b_t^y = \bar{\omega} - \varepsilon - c_t^y = \frac{1}{2}\left(\bar{\omega} - \varepsilon - \frac{\varepsilon}{1+r_t}\right).$$

Asset-market clearing requires private saving to absorb government debt:

$$b_t^y + B_t = 0.$$

Substituting the expression for $b_t^y$, we obtain

$$\frac{1}{2}\left(\bar{\omega} - \varepsilon - \frac{\varepsilon}{1+r_t}\right) + B_t = 0,$$

or equivalently

$$\bar{\omega} - \varepsilon + 2B_t = \frac{\varepsilon}{1+r_t}.$$

Thus the equilibrium gross interest factor is

$$1 + r_t = \frac{\varepsilon}{\bar{\omega} - \varepsilon + 2B_t}.$$

This expression is useful because it shows directly how public debt affects equilibrium returns. Since $B_t < 0$, we have $2B_t < 0$, so the denominator is smaller than in autarky. Therefore

$$1 + r_t > \frac{\varepsilon}{\bar{\omega} - \varepsilon},$$

that is, government borrowing raises the interest rate relative to the autarkic allocation.

At the same time, for the gross interest rate to remain below one we need

$$1 + r_t < 1 \quad \Longleftrightarrow \quad \frac{\varepsilon}{\bar{\omega} - \varepsilon + 2B_t} < 1,$$

which is equivalent to

$$\bar{\omega} - \varepsilon + 2B_t > \varepsilon.$$

Hence a sufficient bound is

$$B_t > \varepsilon - \frac{\bar{\omega}}{2}.$$

The threshold value is therefore

$$B_t = -\frac{\bar{\omega}}{2} + \varepsilon.$$

So, as long as debt is not too negative, we are in the region

$$\frac{\varepsilon}{\bar{\omega} - \varepsilon} < 1 + r_t < 1.$$

This is exactly the interesting case: government debt pushes the interest rate up relative to autarky, but not enough to move the economy out of the dynamically inefficient region.

**Debt dynamics.** With no taxes from period $t \geq 1$ onward, the government rolls over its debt according to

$$B_{t+1} = (1 + r_t)B_t.$$

If $B_t < 0$ and $0 < 1 + r_t < 1$, then

$$B_{t+1} > B_t,$$

so debt becomes less negative over time. In other words, the government keeps rolling over its debt, but because the gross interest factor is below one, the magnitude of debt shrinks over time and converges back toward zero.

As $B_t \to 0$, the equilibrium gross interest factor converges back to its autarkic level

$$1 + r_t \to \frac{\varepsilon}{\bar{\omega} - \varepsilon}.$$

**Welfare intuition.** The initial old are unambiguously better off: they receive the transfer financed by the initial debt issue and, since they are at the end of life, they do not bear the future adjustment. More broadly, the move away from autarky improves the intergenerational allocation because autarky forces a very uneven life-cycle consumption profile, while borrowing allows some partial smoothing. With concave preferences, this smoothing is welfare improving.

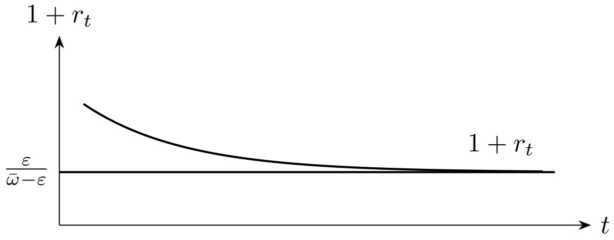

Figure 41: After the initial debt issue, the gross interest factor rises above its autarkic level and then converges back as debt is diluted over time.

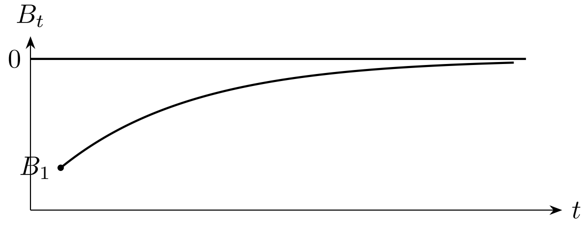

Figure 42: If $B_1 < 0$ and $1 + r_t < 1$, government debt converges upward to zero from below.

# 14 Money Demand

**Reading.** Romer, Section 7.6, pp. 335–338; Sections 6.1–6.3; and Section 7.8.

## 14.1 Why can fiat money have value?

A useful starting point is the Walrasian benchmark. In a complete-markets Arrow–Debreu environment, fiat money is redundant: any desired intertemporal or state-contingent transfer can already be implemented directly through securities trading, so intrinsically useless paper has no essential role.[7]

To generate a demand for money, one must introduce a friction that makes direct credit or complete contingent contracting impossible or too costly. Three canonical routes are:

- **Spatial separation.** If agents are physically separated and do not meet again, private IOUs may fail because repayment cannot be enforced. In that case, fiat money can emerge as a device that carries value across locations and dates.[8]
- **Search and matching frictions.** In decentralized exchange, agents do not meet all potential counterparties simultaneously, and bilateral trade opportunities are limited. A commonly accepted object can then arise endogenously as a medium of exchange.[9]
- **Overlapping generations.** In OLG economies, fiat money can transfer resources across generations in environments where private credit alone cannot implement the relevant allocation.[10]

Thus, the general lesson is simple: *money matters only once we depart from the frictionless Walrasian benchmark.*

## 14.2 Autarky

Consider the two-period endowment profile

$$(\omega_t^y, \omega_{t+1}^o) = (1-\varepsilon, \varepsilon),$$

with logarithmic preferences

$$u(c_t^y, c_{t+1}^o) = \ln c_t^y + \ln c_{t+1}^o.$$

In autarky, agents simply consume their endowments:

$$c_t^y = 1-\varepsilon, \qquad c_{t+1}^o = \varepsilon.$$

Hence the autarkic gross rate of return is pinned down by the Euler equation evaluated at the autarkic allocation:

$$1 + r_t \;=\; \frac{c_{t+1}^o}{c_t^y} \;=\; \frac{\varepsilon}{1-\varepsilon}.$$

This is a useful benchmark: it tells us the intertemporal relative price that is implicit in the no-trade allocation.

## 14.3 Can the government improve on autarky with fiat money?

Yes. The key idea is to endow the initial old with fiat money and ask whether the young are willing to hold it. Let a young agent at time $t$ choose bond holdings $b_t^y$ and nominal money holdings $m_t^y \geq 0$.

---

[7] This is the natural benchmark behind the statement "Walrasian ⇒ no money." For the way modern recursive macro connects this idea to monetary models, see Ljungqvist and Sargent, *Recursive Macroeconomic Theory*, Chapter 28.

[8] The classic reference is Townsend (1980), "Models of Money with Spatially Separated Agents," in *Models of Monetary Economies*, edited by Kareken and Wallace. This is the canonical "turnpike" logic.

[9] See Kiyotaki and Wright (1989), "On Money as a Medium of Exchange," *Journal of Political Economy*, and Kiyotaki and Wright (1993), "A Search-Theoretic Approach to Monetary Economics," *American Economic Review.*

[10] The classic OLG monetary reference is Wallace (1980), "The Overlapping Generations Model of Fiat Money," in *Models of Monetary Economies*, edited by Kareken and Wallace.

Their budget constraints are

$$c_t^y + b_t^y + \frac{m_t^y}{p_t} = \omega_t^y, \tag{14.1}$$

$$c_{t+1}^o - (1 + r_t)b_t^y = \omega_{t+1}^o + \frac{m_t^y}{p_{t+1}}. \tag{14.2}$$

Combining the two constraints yields the present-value budget constraint

$$c_t^y + \frac{c_{t+1}^o}{1 + r_t} = \omega_t^y + \frac{\omega_{t+1}^o}{1 + r_t} + m_t^y \left[\frac{1}{p_{t+1}(1 + r_t)} - \frac{1}{p_t}\right]. \tag{14.3}$$

The associated Lagrangian is

$$\mathcal{L} = \ln c_t^y + \ln c_{t+1}^o + \lambda_t \left\{\omega_t^y + \frac{\omega_{t+1}^o}{1 + r_t} - c_t^y - \frac{c_{t+1}^o}{1 + r_t} + m_t^y \left[\frac{1}{p_{t+1}(1 + r_t)} - \frac{1}{p_t}\right]\right\} + \psi_t m_t^y,$$

where $\psi_t \geq 0$ is the multiplier on the nonnegativity constraint $m_t^y \geq 0$.

The first-order condition with respect to money is

$$\lambda_t \left[\frac{1}{p_{t+1}(1 + r_t)} - \frac{1}{p_t}\right] + \psi_t = 0. \tag{14.4}$$

This condition immediately implies the standard complementary-slackness statement:

$$m_t^y \left[\frac{1}{p_{t+1}(1 + r_t)} - \frac{1}{p_t}\right] = 0. \tag{14.5}$$

Equation (14.5) is the key no-arbitrage restriction. It says:

- if
  $$\frac{p_t}{p_{t+1}} > 1 + r_t,$$
  then money strictly dominates bonds and households would want an arbitrarily large money position;
- if
  $$\frac{p_t}{p_{t+1}} < 1 + r_t,$$
  then money is dominated and optimally
  $$m_t^y = 0;$$
- therefore, in any interior monetary equilibrium,
  $$\frac{p_t}{p_{t+1}} = 1 + r_t. \tag{14.6}$$

So the real return on fiat money must equal the return on bonds along an equilibrium path with valued money.

**Remark.** The mechanism here is not fundamentally an FTPL argument. The essential role of fiat money in this OLG setup is to implement an intergenerational transfer that private markets cannot replicate under autarky. One can certainly reinterpret the initial money injection as a tax-transfer operation financed by seigniorage, but the central logic is not "taxes determine the price level." The central logic is that fiat money relaxes an intergenerational trading restriction.

## 14.4 Household demand for saving and money

With logarithmic preferences, the Euler equation gives

$$\frac{c_{t+1}^{o}}{c_t^{y}} = 1 + r_t. \tag{14.7}$$

Substituting the endowment profile $(1-\varepsilon, \varepsilon)$ into the present-value budget constraint gives

$$c_t^y + \frac{c_{t+1}^o}{1+r_t} = (1-\varepsilon) + \frac{\varepsilon}{1+r_t}.$$

Using (14.7), we obtain

$$c_t^y = \frac{1}{2}\left[1 - \varepsilon + \frac{\varepsilon}{1+r_t}\right]. \tag{14.8}$$

Hence savings by the young are

$$\begin{aligned} s_t^y &= \omega_t^y - c_t^y \\ &= (1-\varepsilon) - \frac{1}{2}\left[1 - \varepsilon + \frac{\varepsilon}{1+r_t}\right] \\ &= \frac{1}{2}\left[1 - \varepsilon - \frac{\varepsilon}{1+r_t}\right]. \end{aligned} \tag{14.9}$$

These savings must be allocated across the two available stores of value:

$$s_t^y = b_t^y + \frac{m_t^y}{p_t}. \tag{14.10}$$

If the nominal money stock supplied by the initial old is $\bar{M}$, money-market clearing implies

$$m_t^y = \bar{M}. \tag{14.11}$$

If, in addition, net bond supply is zero, then bond-market clearing implies

$$b_t^y = 0. \tag{14.12}$$

Combining (14.9), (14.10), (14.11), and (14.12), we obtain

$$\frac{\bar{M}}{p_t} = \frac{1}{2}\left[1 - \varepsilon - \frac{\varepsilon}{1+r_t}\right]. \tag{14.13}$$

Using (14.6), this becomes

$$\frac{\bar{M}}{p_t} = \frac{1}{2}\left[1 - \varepsilon - \varepsilon\frac{p_{t+1}}{p_t}\right]. \tag{14.14}$$

## 14.5 A one-dimensional law of motion

Using the normalization from class, the equilibrium dynamics can be written as

$$P_{t+1} = \frac{\varepsilon}{1-(1-\varepsilon)P_t}. \tag{14.15}$$

A steady state satisfies $P_{t+1} = P_t = P$, hence

$$P = \frac{\varepsilon}{1-(1-\varepsilon)P},$$

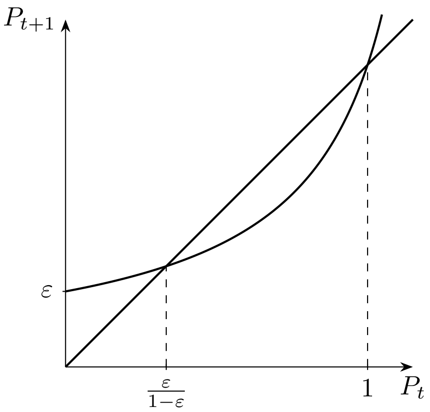

Figure 43: Dynamics implied by (14.15).

or equivalently

$$(1-\varepsilon)P^2 - P + \varepsilon = 0.$$

Therefore the stationary points are

$$P^* = 1 \qquad \text{and} \qquad P^* = \frac{\varepsilon}{1-\varepsilon}. \tag{14.16}$$

Notice that $P^* = 1$ is always a solution, while the second stationary point reproduces the same expression that already appeared in the autarkic gross return.

Figure 14.15 makes the main point visually transparent: the monetary economy may admit multiple stationary equilibria. The shape of the map is economically intuitive. Because the current value of money affects saving decisions, which in turn affect tomorrow's willingness to hold money, the law of motion is nonlinear and can intersect the 45° line more than once.

**Economic interpretation.** Autarky is not generically constrained efficient in the OLG environment. Fiat money can improve on autarky because it creates a store of value that lets a young generation transfer resources to old age. In equilibrium, that same asset is then passed forward to the next generation. The initial old benefit because the government gives them the initial stock of fiat money; the first young are willing to buy it because they expect the next young to do the same. That is the basic recursive logic behind valued fiat money in OLG models.

### 14.5.1 Dynamics of the Price Map and Stability of the Steady States

It is useful to study the law of motion

$$P_{t+1} = F(P_t) := \frac{\varepsilon}{1-(1-\varepsilon)P_t}. \tag{14.17}$$

That is a *nonlinear first-order difference equation*: time is discrete, so the dynamics are generated by iterating the map $F$.

**Step 1: steady states.** A steady state is a fixed point $P^*$ satisfying

$$P^* = F(P^*).$$

Substituting (14.17), we obtain

$$P^* = \frac{\varepsilon}{1-(1-\varepsilon)P^*} \qquad \Longleftrightarrow \qquad (1-\varepsilon)(P^*)^2 - P^* + \varepsilon = 0.$$

Hence the two fixed points are

$$P_L^* = \frac{\varepsilon}{1-\varepsilon}, \qquad P_H^* = 1. \tag{14.18}$$

**Step 2: local stability.** To determine whether these steady states attract nearby paths, we compute the derivative of the map:

$$F'(P) = \frac{\varepsilon(1-\varepsilon)}{[1-(1-\varepsilon)P]^2}.$$

Evaluating at the two fixed points gives

$$F'\left(\frac{\varepsilon}{1-\varepsilon}\right) = \frac{\varepsilon}{1-\varepsilon}, \qquad F'(1) = \frac{1-\varepsilon}{\varepsilon}.$$

Thus, when $\varepsilon < 1/2$,

$$\frac{\varepsilon}{1-\varepsilon} < 1 \qquad \text{and} \qquad \frac{1-\varepsilon}{\varepsilon} > 1.$$

Therefore,

- $P_L^* = \frac{\varepsilon}{1-\varepsilon}$ is *locally stable*;
- $P_H^* = 1$ is *unstable*.

This is already highly informative. The lower fixed point attracts nearby trajectories, whereas the fixed point at 1 is a knife-edge equilibrium: if the economy starts exactly at $P_0 = 1$, it stays there forever, but any small perturbation moves the sequence away from that point.

**Step 3: comparing $F(P_t)$ and $P_t$.** The comparison between $F(P_t)$ and $P_t$ tells us the *direction of motion* of the sequence.

Indeed, since the law of motion is

$$P_{t+1} = F(P_t),$$

we have:

- if $F(P_t) > P_t$, then
$$P_{t+1} > P_t,$$
so the sequence moves *upward* between periods $t$ and $t+1$;
- if $F(P_t) < P_t$, then
$$P_{t+1} < P_t,$$
so the sequence moves *downward*;
- if $F(P_t) = P_t$, then
$$P_{t+1} = P_t,$$
so the economy is already at a fixed point.

This is why the 45°-line diagram is so useful. The graph of $P_{t+1} = F(P_t)$ tells us the next-period value associated with each current value $P_t$, while the line $P_{t+1} = P_t$ identifies the points where the variable does not move. Whenever the graph of $F$ lies *above* the 45° line, we have $F(P_t) > P_t$, so the sequence rises. Whenever it lies *below* the 45° line, we have $F(P_t) < P_t$, so the sequence falls.

**Step 4: sign of $F(P) - P$.** To determine the global direction of motion, consider

$$F(P) - P = \frac{\varepsilon}{1 - (1-\varepsilon)P} - P.$$

A convenient factorization is

$$F(P) - P = \frac{(1-\varepsilon)\left(P - \frac{\varepsilon}{1-\varepsilon}\right)(P-1)}{1-(1-\varepsilon)P}. \tag{14.19}$$

As long as

$$P < \frac{1}{1-\varepsilon},$$

the denominator in (14.19) is positive, so the sign of $F(P) - P$ is determined by the term

$$\left(P - \frac{\varepsilon}{1-\varepsilon}\right)(P-1).$$

Now assume $\varepsilon < 1/2$, so that

$$\frac{\varepsilon}{1-\varepsilon} < 1.$$

Then the dynamics are as follows:

1. If
$$0 < P_t < \frac{\varepsilon}{1-\varepsilon},$$
then both
$$P_t - \frac{\varepsilon}{1-\varepsilon} < 0 \qquad \text{and} \qquad P_t - 1 < 0,$$
so their product is positive. Hence $F(P_t) - P_t > 0$, which implies
$$P_{t+1} > P_t.$$
The sequence moves upward.

2. If
$$\frac{\varepsilon}{1-\varepsilon} < P_t < 1,$$
then
$$P_t - \frac{\varepsilon}{1-\varepsilon} > 0 \qquad \text{but} \qquad P_t - 1 < 0,$$
so the product is negative. Hence $F(P_t) - P_t < 0$, which implies
$$P_{t+1} < P_t.$$
The sequence moves downward.

3. If
$$P_t = 1,$$
then $F(P_t) = P_t$, so the economy remains forever at that steady state. But because the slope there exceeds one, nearby paths move away from it rather than toward it.

4. If
$$1 < P_t < \frac{1}{1-\varepsilon},$$

then both

$$P_t - \frac{\varepsilon}{1-\varepsilon} > 0 \qquad \text{and} \qquad P_t - 1 > 0,$$

so the product is positive. Hence $F(P_t) - P_t > 0$, which implies

$$P_{t+1} > P_t.$$

The sequence again moves upward, now away from the unstable fixed point 1.

**Conclusion.** The lower fixed point

$$P_L^* = \frac{\varepsilon}{1-\varepsilon}$$

is the stable attractor of the system. More precisely:

- if $P_0 \in (0,1)$ and $P_0 \neq 1$, the sequence converges to
$$\frac{\varepsilon}{1-\varepsilon};$$
it rises toward that value if it starts below it, and falls toward that value if it starts above it but still below 1;
- if $P_0 = 1$, the sequence remains forever at 1;
- if $P_0 > 1$, the sequence moves away from 1, so the steady state at 1 is not dynamically relevant except as a knife-edge case.

**Economic interpretation.** The equality $P_H^* = 1$ defines a stationary monetary equilibrium, but it is unstable. By contrast,

$$P_L^* = \frac{\varepsilon}{1-\varepsilon}$$

is dynamically stable and therefore organizes the actual motion of the system. Notice also that this lower fixed point coincides numerically with the gross return associated with autarky. This does *not* mean that the economy literally returns to autarky as an institutional arrangement. Rather, it means that the stable stationary price reproduces the same intertemporal relative price that appeared in the autarkic allocation.

In that sense, the map has a simple economic message: the steady state at 1 exists, but it is fragile; the lower fixed point is the one that attracts nearby economies and therefore governs the observable dynamics of the model.

# 15 Money in the Utility Function

A common reduced-form way to introduce money into an otherwise standard intertemporal model is to assume that households derive utility not only from consumption, but also from holding real balances. Thus, in a deterministic or stochastic environment, preferences can be written as

$$E_0 \sum_{t=0}^{\infty} \beta^t \left[ u(c_t) + \Delta\left(\frac{M_t}{P_t}\right) \right],$$

where $0 < \beta < 1$, $u'(c) > 0$, $u''(c) < 0$, and

$$\Delta'(m) > 0, \qquad \Delta''(m) < 0, \qquad m \equiv \frac{M}{P}.$$

The interpretation is that money provides *liquidity services.* This is analytically convenient, though conceptually somewhat ad hoc: money does not literally enter utility in the same natural way as consumption or leisure. Instead, the specification should be understood as a shortcut for modeling the usefulness of transaction balances.

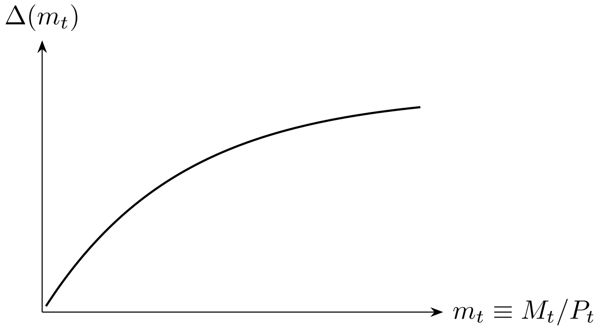

Figure 44: Utility from real balances: increasing and concave.

## 15.1 Nominal Budget Constraint

Let $S_t$ denote one-period nominal bonds purchased in period $t$, paying gross nominal return $1+i_t$ between $t$ and $t+1$. Let $H_t$ denote nominal transfers. A convenient nominal budget constraint is

$$P_t c_t + S_t + M_t = P_t y_t + (1+i_{t-1})S_{t-1} + M_{t-1} + H_t.$$

The household chooses sequences $\{c_t, S_t, M_t\}_{t\geq 0}$ to maximize utility subject to this sequence of constraints and a standard transversality condition.

## 15.2 First-Order Conditions

Using multipliers $\{\lambda_t\}_{t\geq 0}$, the first-order conditions are

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial c_t} &= 0 \quad \Rightarrow \quad u'(c_t) = \lambda_t P_t, \\
\frac{\partial \mathcal{L}}{\partial S_t} &= 0 \quad \Rightarrow \quad \lambda_t = \beta(1+i_t)E_t[\lambda_{t+1}], \\
\frac{\partial \mathcal{L}}{\partial M_t} &= 0 \quad \Rightarrow \quad \lambda_t = \beta E_t[\lambda_{t+1}] + \frac{1}{P_t}\Delta'\left(\frac{M_t}{P_t}\right).
\end{aligned}$$

These conditions have the usual interpretation.

The first condition equates the marginal utility of consumption to the shadow value of nominal resources. The second condition says that the household is indifferent at the optimum between giving up one unit of nominal resources today and saving it in a nominal bond. The third condition says that money is like a bond with a lower pecuniary return, but with an additional liquidity payoff.

## 15.3 Euler Equation

Combining

$$u'(c_t) = \lambda_t P_t \qquad \text{and} \qquad \lambda_t = \beta(1+i_t)E_t[\lambda_{t+1}],$$

we obtain

$$u'(c_t) = \beta(1+i_t)E_t\left[\frac{P_t}{P_{t+1}}u'(c_{t+1})\right].$$

This is the standard intertemporal Euler equation in nominal form. The gross nominal interest rate must compensate the household for postponing consumption, adjusted by expected inflation through the term $P_t/P_{t+1}$.

## 15.4 Money Demand

To derive money demand, subtract the bond FOC from the money FOC:

$$\frac{1}{P_t}\Delta'\left(\frac{M_t}{P_t}\right) = \beta i_t E_t[\lambda_{t+1}].$$

Using

$$\lambda_t = \beta(1+i_t)E_t[\lambda_{t+1}] \quad \Rightarrow \quad \beta E_t[\lambda_{t+1}] = \frac{\lambda_t}{1+i_t},$$

we get

$$\frac{1}{P_t}\Delta'\left(\frac{M_t}{P_t}\right) = \lambda_t\frac{i_t}{1+i_t}.$$

Multiplying by $P_t$ and using $u'(c_t) = \lambda_t P_t$,

$$\Delta'\left(\frac{M_t}{P_t}\right) = u'(c_t)\frac{i_t}{1+i_t}.$$

Hence the household chooses real balances so that the marginal liquidity benefit of money equals its opportunity cost relative to interest-bearing bonds.

A useful implication is immediate: for given consumption, higher nominal interest rates reduce real money demand. Since $\Delta'$ is decreasing, the above condition implies a downward-sloping money demand schedule.

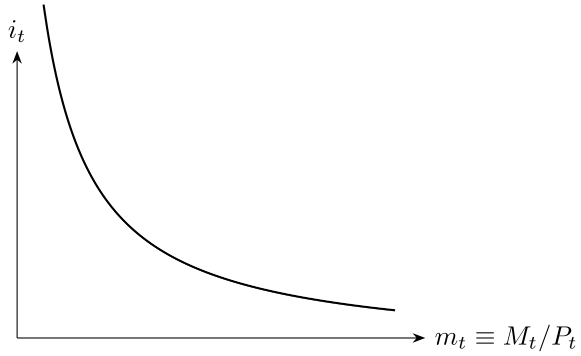

Figure 45: Money demand: higher nominal interest rates imply lower real balances.

## 15.5 Neutrality

In this specification, money typically leaves the *real block* of the model unchanged, provided prices are fully flexible and utility is additively separable in consumption and real balances.

The reason is simple. The Euler equation for consumption and the underlying resource constraints determine real allocations. Once $c_t$, $y_t$, and any other real variables are pinned down, the money-demand equation determines desired real balances $M_t/P_t$. The price level then adjusts to make the nominal quantity of money consistent with that desired level of real balances.

Thus, under flexible prices, changes in the nominal money supply affect nominal variables—most importantly the price level—but do not affect real allocations. This is monetary neutrality.

## 15.6 Friedman Rule

The money-demand condition also delivers the classic Friedman rule. If the policymaker sets the nominal interest rate to zero,

$$i_t = 0,$$

then the opportunity cost of holding money vanishes. Households are no longer penalized for holding transaction balances.

Using the approximate Fisher relation,

$$r_t \approx i_t - \pi_{t+1},$$

the condition $i_t = 0$ implies

$$\pi_{t+1} \approx -r_t.$$

Thus, the Friedman rule prescribes a mild deflation at roughly the real interest rate, so that the nominal return on money matches that on bonds.

Intuitively, money is socially cheap to produce, but privately costly to hold whenever $i_t > 0$. Optimal policy eliminates this wedge by driving the nominal interest rate to zero.

## 15.7 Why Is Money Often Non-Neutral in Practice?

The neutrality result is powerful, but it depends on strong assumptions. In the data, monetary policy appears to have real effects, at least in the short run. The usual explanation is nominal rigidity.

If prices or wages adjust slowly, then a monetary disturbance changes real balances, real interest rates, and aggregate demand before the price level fully responds. Under those conditions, monetary policy affects output and employment temporarily, so money is not neutral in the short run.

The main mechanisms emphasized in the literature are:

- sticky prices,
- sticky wages,
- menu costs,
- staggered adjustment,
- coordination failures in price-setting.

This is the basic logic behind New Keynesian models. In that class of environments, the MIU specification can still be used to model money demand, but the key source of non-neutrality comes from price adjustment frictions rather than from money directly entering utility.

Empirically, the broad evidence is consistent with short-run monetary non-neutrality and incomplete nominal adjustment; see, for example, Christiano, Eichenbaum, and Evans (2005) and Nakamura and Steinsson (2008).

# 16 Evidence on Price Rigidity

A useful survey is Klenow and Malin (2010), who document that many prices adjust infrequently, though the frequency of adjustment varies substantially across sectors and goods. This empirical evidence motivates theories in which firms do not continuously reset prices.

# 17 Sticky Prices and Menu Costs

A simple way to rationalize nominal rigidity is to assume that a firm must pay a fixed cost $z > 0$ whenever it changes its nominal price. This is the classic *menu cost* idea associated with Mankiw (1985) and Akerlof and Yellen (1985).

## 17.1 Firm's problem

Let the firm's real profit be

$$\Pi\left(\frac{P_i}{P}\right),$$

so profits depend on the firm's *relative price.* If all firms are symmetric, then in a symmetric equilibrium $P_i = P$, and it is convenient to think of profits as a function of a single real price variable

$$p \equiv \frac{P_i}{P}.$$

Let $p^*$ denote the firm's privately optimal relative price. Then

$$\Pi'(p^*) = 0, \qquad \Pi''(p^*) < 0.$$

The second condition says that profits are locally concave around the optimum.

## 17.2 The private loss from not adjusting

Suppose the firm's current price is $p \neq p^*$, but changing the price requires paying the fixed cost $z$. Define the private loss from not adjusting as

$$L(p) \equiv \Pi(p^*) - \Pi(p) \geq 0.$$

A second-order Taylor expansion of profits around $p^*$ gives

$$\Pi(p) \approx \Pi(p^*) + \Pi'(p^*)(p - p^*) + \frac{1}{2}\Pi''(p^*)(p - p^*)^2.$$

Since $\Pi'(p^*) = 0$,

$$\Pi(p) \approx \Pi(p^*) + \frac{1}{2}\Pi''(p^*)(p - p^*)^2,$$

and therefore

$$L(p) = \Pi(p^*) - \Pi(p) \approx -\frac{1}{2}\Pi''(p^*)(p - p^*)^2.$$

Because $\Pi''(p^*) < 0$, this expression is positive.

This is the key point: *near the optimum, the firm's private loss from keeping the wrong price is quadratic in the deviation* $p - p^*$. Hence the loss is *second order.*

## 17.3 Adjustment rule and inaction region

The firm adjusts if the gain from resetting exceeds the menu cost:

$$L(p) \geq z.$$

Using the quadratic approximation,

$$-\frac{1}{2}\Pi''(p^*)(p - p^*)^2 \geq z.$$

Equivalently,

$$(p - p^*)^2 \geq \frac{2z}{-\Pi''(p^*)}.$$

Thus the firm does *not* adjust whenever the current price lies in a band around the optimum:

$$|p - p^*| < \sqrt{\frac{2z}{-\Pi''(p^*)}}.$$

Define the two thresholds

$$\hat{p}^L = p^* - \sqrt{\frac{2z}{-\Pi''(p^*)}}, \qquad \hat{p}^H = p^* + \sqrt{\frac{2z}{-\Pi''(p^*)}}.$$

Then the firm's policy is:

$$\text{do not adjust if } p \in [\hat{p}^L, \hat{p}^H], \qquad \text{adjust if } p \notin [\hat{p}^L, \hat{p}^H].$$

So even a small fixed adjustment cost can generate a nontrivial *region of inaction.*

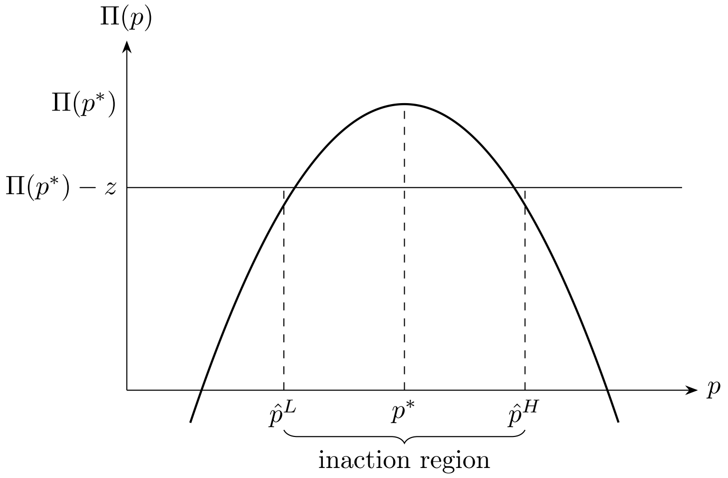

Figure 46: The firm adjusts only when the private gain from resetting exceeds the menu cost $z$.

## 17.4 Economic interpretation

The graph above should be read very carefully.

At $p^*$, profits are maximized, so the slope is zero:

$$\Pi'(p^*) = 0.$$

That is exactly why the loss from a small deviation is second order rather than first order. If the firm is *slightly* away from its optimum, the private cost of doing nothing is tiny.

This is the logic behind the menu cost argument:

> small physical costs of changing prices can rationalize substantial nominal rigidity, because the firm's private loss from staying slightly off its optimum is very small.

## 17.5 Why this does *not* mean rigidity is socially harmless

A very common mistake is to jump from

private loss is second order

to

social loss is also second order.

That conclusion is generally false.

The reason is that the firm's objective and society's objective are not the same. Under monopolistic competition, the firm's privately optimal price is a monopoly price, not the socially efficient price.

Let $p^m$ denote the monopoly price and $p^s$ the socially efficient price, with

$$p^m > p^s.$$

Private profits are maximized at $p^m$, but social surplus is maximized at $p^s$.

Thus, around $p^m$,

$$\Pi'(p^m) = 0, \qquad S'(p^m) \neq 0,$$

where $S(p)$ denotes social surplus.

So a small change in price around $p^m$ has:

- a *second-order* effect on the firm's private profits,
- but a *first-order* effect on social welfare.

This is the precise sense in which price rigidity can be privately cheap but socially costly.

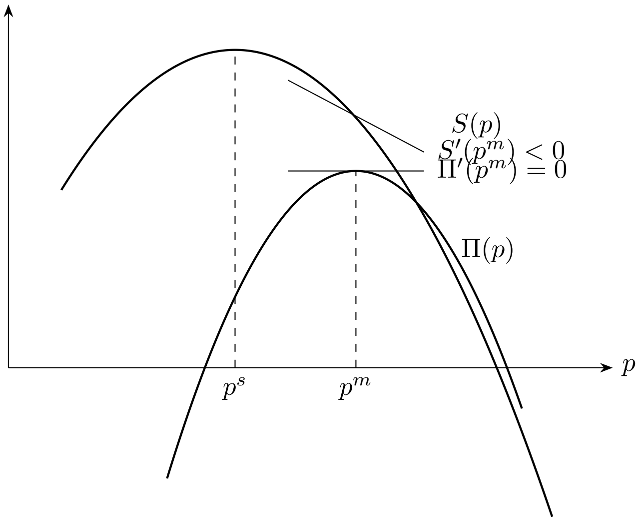

Figure 47: At the monopoly price $p^m$, private profits are locally flat, but social surplus is not.

### 17.6 Bottom line

The menu cost logic has two distinct steps.

**Step 1: private incentives.** Because $\Pi'(p^*) = 0$, the firm's private loss from a small pricing mistake is second order:

$$L(p) \approx -\frac{1}{2}\Pi''(p^*)(p - p^*)^2.$$

Hence even a small fixed adjustment cost can induce the firm not to change its price.

**Step 2: aggregate consequences.** Because the private optimum does not generally coincide with the social optimum, the welfare effects of a firm's pricing decision need not be second order. With monopolistic competition and aggregate-demand externalities, small private frictions can generate large real effects.

This is the central insight of the menu cost literature:

> *small costs of nominal adjustment can produce substantial nominal rigidity and sizable aggregate effects.*

# 18 Dynamic Price Setting

We now move from the static monopoly pricing problem to a dynamic environment in which firms may be unable to reset prices every period.

### 18.1 Static desired markup

Consider a monopolistically competitive firm $i$ facing the isoelastic demand curve

$$Y_t(i) = \left(\frac{P_t(i)}{P_t}\right)^{-\varepsilon} Y_t, \qquad \varepsilon > 1.$$

Suppose the firm's nominal marginal cost is $MC_t$. Then nominal profits are

$$\Pi_t(i) = \big(P_t(i) - MC_t\big)Y_t(i).$$

Substituting demand,

$$\Pi_t(i) = \big(P_t(i) - MC_t\big)\left(\frac{P_t(i)}{P_t}\right)^{-\varepsilon} Y_t = P_t^{\varepsilon} Y_t \left[P_t(i)^{1-\varepsilon} - MC_t P_t(i)^{-\varepsilon}\right].$$

Differentiating with respect to $P_t(i)$,

$$\frac{\partial \Pi_t(i)}{\partial P_t(i)} = P_t^{\varepsilon} Y_t \left[(1-\varepsilon)P_t(i)^{-\varepsilon} + \varepsilon MC_t P_t(i)^{-\varepsilon-1}\right].$$

Setting this equal to zero and multiplying through by $P_t(i)^{\varepsilon+1}$,

$$(1-\varepsilon)P_t(i) + \varepsilon MC_t = 0.$$

Hence the optimal price is

$$P_t^*(i) = \frac{\varepsilon}{\varepsilon - 1} MC_t.$$

Define the desired markup

$$\mu \equiv \frac{\varepsilon}{\varepsilon - 1} > 1.$$

Then the firm's desired price is simply

$$P_t^*(i) = \mu \, MC_t.$$

Thus, under flexible prices, the monopolistically competitive firm sets price as a constant markup over nominal marginal cost.

**Empirical benchmark.** If the average markup is about 10%, then $\mu = 1.1$, so

$$1.1 = \frac{\varepsilon}{\varepsilon - 1} \qquad \Longrightarrow \qquad \varepsilon = 11.$$

This is a useful quantitative benchmark.

## 18.2 Real marginal cost and log-linear form

Let

$$mc_t \equiv \frac{MC_t}{P_t}$$

denote real marginal cost. Dividing the pricing rule by the aggregate price level,

$$\frac{P_t^*(i)}{P_t} = \mu \, mc_t.$$

If we log-linearize around a zero-inflation steady state with constant desired markup, then in hat notation the optimal reset price satisfies

$$\hat{p}_t^* = \widehat{mc}_t.$$

So, in log deviations, the desired relative price moves one-for-one with real marginal cost.

## 18.3 Why pricing becomes dynamic

The static formula above is appropriate only if firms can freely reoptimize every period. In practice, this may fail for several reasons:

- fixed menu costs,
- random adjustment opportunities as in Calvo pricing,
- convex costs of price adjustment as in Rotemberg pricing.

In all these cases, the firm must think dynamically: if it resets its price today, that price may remain in place for several periods, so the firm cares not only about current marginal cost, but also about expected future marginal cost.

## 18.4 Calvo pricing

Assume now that in each period a firm can reset its price with probability $1 - \theta$, while with probability $\theta$ it must keep its previous price. This is the Calvo assumption.

A firm that gets to reset in period $t$ chooses $P_t^*$ to maximize the expected discounted value of profits over all future periods in which that price may still be in effect:

$$\max_{P_t^*} E_t \sum_{k=0}^{\infty} (\beta\theta)^k \Lambda_{t,t+k} \Big( P_t^* - MC_{t+k} \Big) Y_{t+k|t},$$

subject to

$$Y_{t+k|t} = \left(\frac{P_t^*}{P_{t+k}}\right)^{-\varepsilon} Y_{t+k}.$$

Here:

- $\beta$ is the household discount factor,
- $\theta$ is the probability the price remains unchanged,
- $\Lambda_{t,t+k}$ is the stochastic discount factor,
- $Y_{t+k|t}$ is demand in period $t+k$ for a firm that last reset in period $t$.

The first-order condition can be written as

$$E_t \sum_{k=0}^{\infty} (\beta\theta)^k \Lambda_{t,t+k} Y_{t+k|t} \Big(P_t^* - \mu MC_{t+k}\Big) = 0.$$

This makes the economic logic transparent:

> the optimal reset price is a weighted average of current and expected future desired markup prices.

If the firm expects future marginal costs to be high, it will choose a higher reset price today.

## 18.5 Log-linearized reset price equation

After log-linearization around a zero-inflation steady state, the reset price satisfies

$$\hat{p}_t^* = (1-\beta\theta)\widehat{mc}_t + \beta\theta E_t\left[\hat{p}_{t+1}^* + \pi_{t+1}\right].$$

This equation is central. It says:

- the reset price is higher when current real marginal cost is high;
- the reset price is also higher when firms expect future reset prices and future inflation to be high.

So pricing becomes forward-looking.

## 18.6 From reset prices to inflation

Under Calvo pricing, the aggregate price level evolves according to

$$P_t^{1-\varepsilon} = \theta P_{t-1}^{1-\varepsilon} + (1-\theta)(P_t^*)^{1-\varepsilon}.$$

Log-linearizing this around the zero-inflation steady state yields

$$\pi_t = \frac{1-\theta}{\theta}\hat{p}_t^*.$$

Combining this with the reset-price equation gives the New Keynesian Phillips Curve:

$$\pi_t = \lambda\widehat{mc}_t + \beta E_t[\pi_{t+1}],$$

where

$$\lambda \equiv \frac{(1-\theta)(1-\beta\theta)}{\theta}.$$

This is the fundamental inflation equation in the baseline New Keynesian model.

**Interpretation.** Inflation is high when:

- current real marginal cost is high;
- firms expect future inflation to be high.

Thus inflation is both *cost-driven* and *forward-looking.*

### 18.7 Rotemberg pricing

An alternative to Calvo is to assume convex costs of price adjustment, as in Rotemberg pricing. Instead of some firms being unable to adjust, all firms can adjust every period, but changing prices is costly. After log-linearization, this approach delivers an inflation equation very similar to the Calvo-based New Keynesian Phillips Curve. For many purposes, the two frameworks lead to comparable aggregate implications.

## 19 Marginal Cost in the Baseline New Keynesian Model

To obtain an output-gap representation of inflation, we now express real marginal cost in terms of output and productivity.

### 19.1 Labor supply

Assume labor supply takes the form

$$\frac{W_t}{P_t} = \psi C_t^{1/\sigma} N_t^{1/\eta},$$

where:

- $\sigma > 0$ governs intertemporal substitution in consumption,
- $\eta > 0$ governs the Frisch elasticity of labor supply,
- $\psi > 0$ is a scale parameter.

### 19.2 Technology

Assume a one-good economy with linear production in labor:

$$Y_t = Z_t N_t.$$

Since there is only one good,

$$Y_t = C_t.$$

Nominal marginal cost is the nominal wage divided by productivity,

$$MC_t = \frac{W_t}{Z_t}.$$

Therefore real marginal cost is

$$mc_t = \frac{MC_t}{P_t} = \frac{W_t}{P_t Z_t}.$$

Using labor supply and $C_t = Y_t$,

$$mc_t = \frac{1}{Z_t}\frac{W_t}{P_t} = \frac{1}{Z_t}\psi C_t^{1/\sigma} N_t^{1/\eta} = \frac{1}{Z_t}\psi Y_t^{1/\sigma}\left(\frac{Y_t}{Z_t}\right)^{1/\eta}.$$

Thus

$$mc_t = \psi Y_t^{1/\sigma+1/\eta} Z_t^{-(1+1/\eta)}.$$

Taking logs and log-linearizing,

$$\widehat{mc}_t = \left(\frac{1}{\eta} + \frac{1}{\sigma}\right) \hat{y}_t - \left(\frac{\eta+1}{\eta}\right) \hat{z}_t.$$

## 19.3 Flexible-price output

Under perfectly flexible prices, firms always set their desired markup, so real marginal cost is constant:

$$\widehat{mc}_t^{flex} = 0.$$

Hence flexible-price output $\hat{y}_t^{flex}$ satisfies

$$0 = \left(\frac{1}{\eta} + \frac{1}{\sigma}\right) \hat{y}_t^{flex} - \left(\frac{\eta+1}{\eta}\right) \hat{z}_t.$$

Subtracting this relation from the expression for $\widehat{mc}_t$ gives

$$\widehat{mc}_t = \left(\frac{1}{\eta} + \frac{1}{\sigma}\right) \left(\hat{y}_t - \hat{y}_t^{flex}\right).$$

Thus real marginal cost is proportional to the *output gap*.

## 19.4 Phillips curve in output-gap form

Substituting into the New Keynesian Phillips Curve,

$$\pi_t = \lambda \left(\frac{1}{\eta} + \frac{1}{\sigma}\right) \left(\hat{y}_t - \hat{y}_t^{flex}\right) + \beta E_t[\pi_{t+1}].$$

Defining the output gap

$$x_t \equiv \hat{y}_t - \hat{y}_t^{flex},$$

we can write

$$\pi_t = \kappa x_t + \beta E_t[\pi_{t+1}],$$

where

$$\kappa = \lambda \left(\frac{1}{\eta} + \frac{1}{\sigma}\right).$$

This is the standard New Keynesian Phillips Curve in output-gap form.

**Economic interpretation.** Inflation rises when output is above its flexible-price level, because production above the natural allocation requires higher labor input, raises wages, raises marginal cost, and therefore induces firms to set higher prices.

**Summary.** The full logic is:

output gap $\implies$ higher real marginal cost $\implies$ higher desired reset prices $\implies$ higher inflation.

# 20 The Simple New Keynesian Model

We can summarize the simple New Keynesian model with three equations:

$$\tilde{\pi}_t = \lambda \left(\frac{1}{\eta} + \frac{1}{\sigma}\right) \hat{Y}_t + \beta E_t[\tilde{\pi}_{t+1}], \tag{20.1}$$

$$\hat{Y}_t = E_t\left[\hat{Y}_{t+1} - \sigma\beta i_t + \sigma\tilde{\pi}_{t+1} + \varepsilon_t^{IS}\right], \tag{20.2}$$

$$i_t = \phi_0 + \phi_\pi \tilde{\pi}_t + \phi_Y \hat{Y}_t + \varepsilon_t^i. \tag{20.3}$$

Equation (20.1) is the New Keynesian Phillips Curve (NKPC), equation (20.2) is the forward-looking IS curve, and equation (20.3) is the Taylor rule.

## 20.1 A useful graphical interpretation

To build intuition, it is useful to treat expected future variables as given in the short run. That is, for the purpose of the graph, take $E_t[\hat{Y}_{t+1}]$ and $E_t[\tilde{\pi}_{t+1}]$ as fixed.

**Step 1: IS and Taylor rule in the $(\hat{Y}_t, i_t)$-plane.** From the IS curve,

$$\hat{Y}_t = A_t - \sigma\beta i_t,$$

where

$$A_t \equiv E_t[\hat{Y}_{t+1}] + \sigma E_t[\tilde{\pi}_{t+1}] + \varepsilon_t^{IS}.$$

Thus the IS curve is downward sloping in the $(\hat{Y}_t, i_t)$-plane:

$$i_t = \frac{A_t - \hat{Y}_t}{\sigma\beta}.$$

For a given inflation rate $\tilde{\pi}_t$, the Taylor rule is

$$i_t = \underbrace{\phi_0 + \phi_\pi \tilde{\pi}_t + \varepsilon_t^i}_{\text{intercept}} + \phi_Y \hat{Y}_t,$$

so it is upward sloping in the $(\hat{Y}_t, i_t)$-plane.

A negative monetary policy shock, $\varepsilon_t^i < 0$, shifts the Taylor rule downward. The new intersection features:

$$\hat{Y}_t \uparrow, \qquad i_t \downarrow .$$

**Step 2: NKPC in the $(\hat{Y}_t, \tilde{\pi}_t)$-plane.** The NKPC can be written as

$$\tilde{\pi}_t = \underbrace{\beta E_t[\tilde{\pi}_{t+1}]}_{\text{intercept}} + \underbrace{\lambda \left(\frac{1}{\eta} + \frac{1}{\sigma}\right)}_{\kappa > 0} \hat{Y}_t.$$

Hence the NKPC is upward sloping in the $(\hat{Y}_t, \tilde{\pi}_t)$-plane.

So once the monetary easing raises output from $\hat{Y}_t$ to $\hat{Y}_t'$, inflation rises along the NKPC:

$$\hat{Y}_t \uparrow \implies \tilde{\pi}_t \uparrow .$$

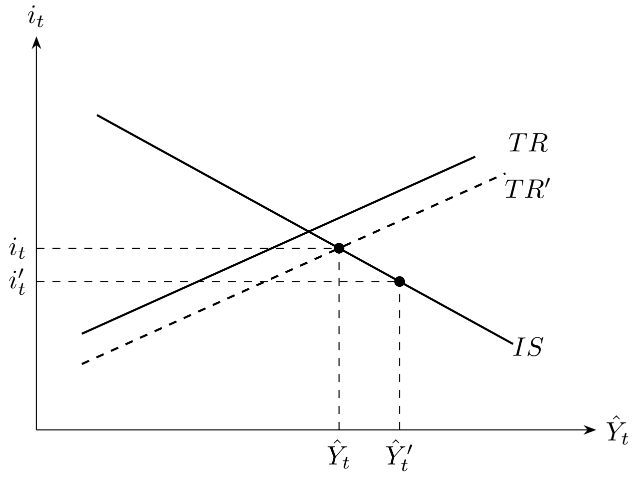

Figure 48: A negative monetary policy shock shifts the Taylor rule downward: output rises and the nominal interest rate falls on impact.

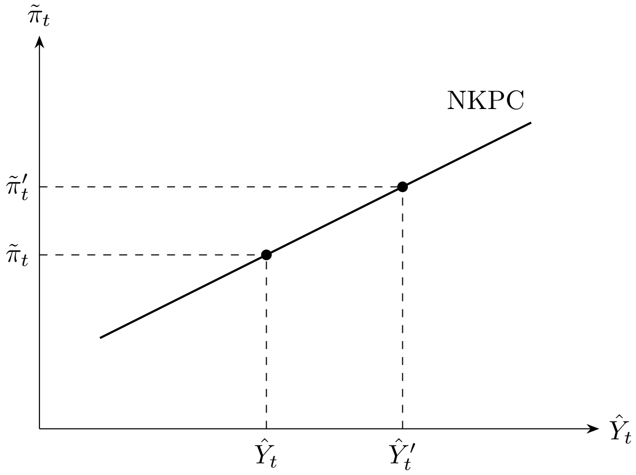

Figure 49: Higher output implies higher inflation through the New Keynesian Phillips Curve.

## 20.2 Putting the two graphs together

The two graphs should be read sequentially:

1. A negative monetary policy shock shifts the Taylor rule downward.
2. This lowers the nominal interest rate and raises output.
3. Higher output increases marginal cost.
4. Through the NKPC, higher marginal cost raises inflation.

So, for an expansionary monetary policy shock, the short-run comparative statics are

$$\hat{Y}_t \uparrow, \qquad i_t \downarrow, \qquad \tilde{\pi}_t \uparrow .$$

## 20.3 What about the "second-round effects"?

At first sight, one may think there is an endless feedback loop:

> lower nominal interest rate → higher output → higher inflation → lower real interest rate → even higher output → $\cdots$

This is a very good question, but the answer is that the model is solved *simultaneously*, not by mechanically iterating causal arrows forever.

The key object is the ex ante real rate,

$$r_t^{real} \equiv i_t - E_t[\tilde{\pi}_{t+1}],$$

up to notation and linearization constants. A monetary easing lowers the nominal rate $i_t$, which tends to lower the real rate and stimulate demand. But as output and inflation rise, the Taylor rule implies that the central bank reacts:

$$i_t = \phi_0 + \phi_\pi \tilde{\pi}_t + \phi_Y \hat{Y}_t + \varepsilon_t^i.$$

So the nominal rate is not fixed forever. Rising inflation and rising output push it back up.

This is exactly why the Taylor principle matters. If

$$\phi_\pi > 1,$$

then an increase in inflation leads to a more than one-for-one increase in the nominal interest rate, so the real interest rate rises rather than falls. That stabilizes the system and rules out explosive second-round dynamics.

## 20.4 A useful distinction

It is important to distinguish different shocks:

**Monetary policy shock $\varepsilon_t^i < 0$.** This shifts the Taylor rule downward. Then it is possible to have

$$\hat{Y}_t \uparrow, \qquad \tilde{\pi}_t \uparrow, \qquad i_t \downarrow .$$

**Positive IS shock $\varepsilon_t^{IS} > 0$.** This shifts the IS curve outward. Then output and inflation rise, but now the Taylor rule typically implies

$$i_t \uparrow .$$

So the statement "output rises, inflation rises, and the nominal rate falls" is *not* a generic property of the model. It is the comparative static associated with an *expansionary monetary policy shock.*

## 20.5 A reduced-form AD relation

Substituting the Taylor rule into the IS curve gives

$$\hat{Y}_t = E_t[\hat{Y}_{t+1} + \sigma\tilde{\pi}_{t+1} + \varepsilon_t^{IS}] - \sigma\beta\big(\phi_0 + \phi_\pi\tilde{\pi}_t + \phi_Y\hat{Y}_t + \varepsilon_t^i\big).$$

Rearranging,

$$(1 + \sigma\beta\phi_Y)\hat{Y}_t = A_t - \sigma\beta\phi_\pi\tilde{\pi}_t,$$

for a suitable constant term $A_t$. Thus

$$\hat{Y}_t = \frac{A_t}{1 + \sigma\beta\phi_Y} - \frac{\sigma\beta\phi_\pi}{1 + \sigma\beta\phi_Y}\tilde{\pi}_t.$$

This is a downward-sloping aggregate demand relation in the $(\hat{Y}_t, \tilde{\pi}_t)$-plane.

Combining this downward-sloping AD curve with the upward-sloping NKPC gives the equilibrium pair $(\hat{Y}_t, \tilde{\pi}_t)$.

### 20.6 Bottom line

The clean intuition is the following:

- the IS curve links demand to the real interest rate;
- the Taylor rule determines how the nominal rate reacts;
- the NKPC links output to inflation;
- together, the three equations pin down output, inflation, and the nominal interest rate.

For a monetary easing:

$$i_t \downarrow \Longrightarrow \hat{Y}_t \uparrow \Longrightarrow \tilde{\pi}_t \uparrow,$$

but the increase in inflation feeds back into policy through the Taylor rule. There is no infinite causal chain: equilibrium is determined jointly, and under the Taylor principle the system is stable.

## 21 Basic Facts about Price Rigidity

A useful starting point for New Keynesian economics is the empirical observation that many individual prices do not adjust continuously. At the micro level, prices often remain fixed for long periods and then change discretely in relatively large jumps.

### 21.1 Why this matters

In a frictionless environment, firms would continuously reset prices in response to changes in marginal cost, demand, or aggregate conditions. The data look very different. Individual prices often display *spells of inaction* followed by occasional adjustments. This is the basic empirical motivation for models of sticky prices.

### 21.2 Core empirical facts

The main stylized facts emphasized in class are the following:

1. **Prices are sticky.** Many individual prices remain unchanged for substantial periods of time.
2. **Price adjustment is lumpy.** When prices change, they typically move in discrete jumps rather than in a smooth continuous way.
3. **Observed price changes are often large.** Conditional on changing, prices usually move by amounts that are economically meaningful, not by infinitesimal margins.
4. **Temporary sales matter.** A large fraction of observed price changes reflect temporary promotions rather than permanent changes in regular prices.
5. **Regular prices are much stickier than posted prices.** Once temporary sales are stripped out, the underlying regular price typically changes much less often.

6. **There is substantial heterogeneity across goods.** Some prices change frequently, while others remain fixed for long periods.
7. **Nominal prices often follow a step pattern.** In low-frequency data, many prices look like a staircase: long flat segments followed by infrequent upward revisions.

## 21.3 Evidence discussed in class

The evidence presented in class can be read as illustrating these facts in a very direct way.

**Example 1: a grocery product across stores.** For Fleischmann's margarine, the price in one store remains nearly fixed for long stretches even while the average price in other stores moves over time. This suggests that firms do not continuously track changing market conditions with immediate price adjustment.

At the same time, the series displays a few sharp temporary spikes and drops. These movements look more like short-lived sales or promotions than permanent changes in the underlying regular price.

**Example 2: posted prices versus regular prices.** For products such as Coca-Cola, Nescafé, bottled water, and deodorant, the posted price fluctuates much more than the regular price. This distinction is crucial.

If one looks only at posted prices, one may conclude that prices are highly flexible. But much of that apparent flexibility reflects temporary discounts. The regular price is typically much smoother and more persistent, which is much more consistent with sticky-price models.

**Example 3: newspaper cover prices.** Newspaper prices provide a very clean illustration of nominal rigidity. In levels, nominal prices remain constant for long periods and then jump discretely. In logs, the same pattern looks like an upward staircase.

This is exactly what one would expect if firms tolerate a gradual erosion of their relative price due to trend inflation and then occasionally reset to a new higher level.

## 21.4 Interpretation

These facts are important because they discipline theory.

- The fact that prices remain unchanged for long periods motivates *price rigidity*.
- The fact that changes occur in jumps motivates models with *fixed adjustment costs* or *infrequent reoptimization*.
- The importance of temporary sales suggests that one must distinguish between the *posted price* and the *regular price*.
- The heterogeneity across goods suggests that no single frictionless benchmark can explain all observed pricing patterns.

Thus, the evidence does not simply say that "prices move slowly." It says something more precise: many firms follow a pricing pattern characterized by *inaction most of the time and occasional discrete adjustments*.

## 21.5 Connection with New Keynesian models

These empirical patterns motivate the main pricing frictions used in modern macroeconomics:

1. **Menu cost models:** firms can adjust whenever they want, but changing prices is costly.
2. **Calvo models:** firms adjust only when they receive a random opportunity.
3. **Rotemberg models:** price adjustment is always possible, but large changes are costly.

Each of these frameworks tries to capture, in a different way, the same broad empirical message from the micro data: prices are not continuously and costlessly adjusted.

## 21.6 A compact summary to remember

A good one-line summary is:

> *Micro prices are not smooth. They are typically sticky, lumpy, and heavily affected by temporary sales.*

A slightly longer summary is:

> *The data show long spells of unchanged regular prices, occasional large adjustments, and many transitory deviations due to promotions. This is the empirical foundation for sticky-price models.*