Lecture 2: Building the New Keynesian Model

From households and price-setting firms to three equations

Natassa Zervou-Velez

International Monetary Relations • AUEB • Winter Semester 2026–27

Outline

Why build the model

Households: the demand side

Firms: price setting and sticky prices

From marginal cost to the output gap

How the model works

Motivation

From last week to today

- Last week: the Taylor rule as a description of what the Fed does: the policy rate rises when inflation is above target and when output is above potential.
- This lecture and the next: why would a central bank behave this way? Is it a good way to set rates?
- To answer, we need a model in which
  - monetary policy matters for output and inflation, and
  - we can ask what policy is optimal, i.e. best for households.

The plan for two lectures

1. Today: build the New Keynesian (NK) model from the decisions of households and firms, and derive its three equations.
2. Next: optimal monetary policy in that model — and show that it looks like a Taylor rule.

Why sticky prices matter

Benchmark: flexible prices

- Firms take prices as given (perfect competition).
- Prices adjust immediately to any shock.
- A change in i_t is matched by a change in expected inflation: the real rate does not move.
- Money is neutral: monetary policy does not move output.

New Keynesian

- Firms have market power (monopolistic competition).
- Firms set their nominal prices.
- Changing prices is costly: prices are sticky.
- Monetary policy moves the real interest rate and output in the short run.

Three steps

(1) Firms as price setters → (2) add costs of changing prices → (3) study optimal monetary policy.

The economy

- Households: consume, work, and save in a nominal bond. "Cashless": no money in the utility function or the budget constraint.
- Firms, in two layers:
  - Retail firms: perfectly competitive; bundle many differentiated goods into one final good that households buy.
  - Wholesale firms: a continuum indexed by $i \in [0, 1]$; each produces a differentiated good, sets its price, and pays a cost to change it.
- Central bank: sets the nominal interest rate i_t.

Why cashless?

In modern economies the central bank sets an interest rate, not the money supply. Money exists in the background, but policy works through i_t.

Roadmap: what we will do today

1. Households. Solve their problem ⇒ two conditions:
   - how much to consume today versus tomorrow (Euler equation);
   - how much to work (labor supply).
2. Firms. Solve their problem ⇒
   - the demand for each good, and the markup under flexible prices;
   - with costs of changing prices: a Phillips curve linking inflation to marginal cost.
3. Combine. Express marginal cost in terms of the output gap.
4. Result: three equations — IS curve, Phillips curve, policy rule.

How to read the derivations

Each step is: write the problem → take first-order conditions → interpret → approximate around the steady state. The intuition matters more than the algebra.

Toolkit: notation

- Upper case for levels, lower case for logs: $c_t \equiv \ln C_t$, $y_t \equiv \ln Y_t$, $n_t \equiv \ln N_t$.
- Steady state: the point where all variables are constant and there are no shocks. A bar denotes a steady-state value: $\bar{c}$, $\bar{n}$.
- The difference of logs is a percentage deviation: $c_t - \bar{c} = \ln(C_t/\bar{C}) \approx (C_t - \bar{C})/\bar{C}$.
- Once we linearize, lower-case variables will denote these percentage deviations from steady state. We use a hat only for marginal cost, $\widehat{mc}_t$.
- Rates (interest, inflation) stay as rates: i_t, π_t.
- E_t: expectation using the information available in period t.

Households

The household problem

$$\max_{\{C_t, N_t, B_t\}} \; E_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{C_t^{1-\sigma}}{1-\sigma} - \chi \frac{N_t^{1+\varphi}}{1+\varphi} \right]$$

subject to the budget constraint in every period:

$$P_t C_t + B_t = (1 + i_{t-1}) B_{t-1} + W_t N_t + D_t$$

- C_t: consumption of the final good; N_t: hours worked.
- B_t: one-period nominal bond, paying the nominal rate i_t.
- W_t: nominal wage; D_t: dividends — firms' profits, paid to households, who own the firms.
- σ: inverse of the intertemporal elasticity of substitution; φ: inverse Frisch elasticity of labor supply.

The functional form (constant relative risk aversion in consumption, constant Frisch elasticity of labor supply) makes linearization transparent.

What the household chooses, and what it takes as given

Chosen by the household

- Consumption C_t
- Hours worked N_t
- Savings in bonds B_t

Taken as given

- Prices: P_t, W_t, i_t
- Dividends D_t
- Past savings B_{t-1}

- Each household is small: it cannot affect prices, wages or the interest rate.
- Prices and the interest rate are determined in equilibrium (by firms and by the central bank).
- D_t is received as a lump sum: it does not depend on the household's own choices, so there is no first-order condition for it.
- Timing: a bond bought in t pays $(1 + i_t)$ in $t + 1$. The rate is known when the bond is bought.

Reading the household problem

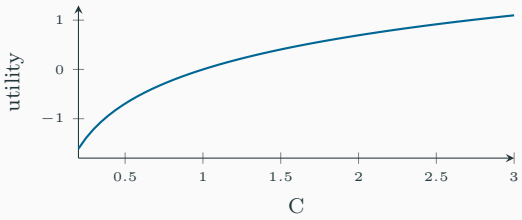
- Objective: the household values consumption and dislikes work, today and in all future periods.
  - $\beta \in (0, 1)$: the discount factor. Future utility counts less than present utility. Quarterly, $\beta \approx 0.99$.
- Budget constraint, in euros:
  - Uses of funds (left): spending on goods $P_t C_t$ and new savings in bonds B_t.
  - Sources of funds (right): bonds bought last period plus interest, labor income, and dividends.
- Choices: each period, how much to consume, how much to work, and how much to save.

Two trade-offs

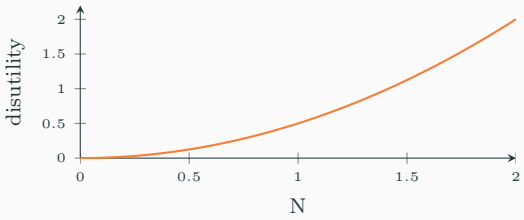
Consume or save? (today versus tomorrow) Work or rest? (consumption versus leisure)

Concave utility, convex disutility of work

Utility of consumption: concave



Disutility of work: convex



- Marginal utility $C_t^{-\sigma}$ falls with C_t: each extra unit of consumption adds less utility.
- Marginal disutility χN_t^{φ} rises with N_t: each extra hour is more costly.
- Why $1-\sigma$ and $1+\varphi$? Dividing by the exponent keeps the derivatives simple: $C_t^{-\sigma}$ and χN_t^{φ}.
- $\sigma = 1$: $\ln C_t$, also concave. Larger σ: more concave, stronger desire to smooth. $\sigma = 0$: linear, no desire to smooth.

Figures: $\ln C$ ($\sigma = 1$) and $N^2/2$ ($\chi = 1$, $\varphi = 1$).

What do σ and φ mean?

σ: willingness to substitute over time

- $1/\sigma$ is the intertemporal elasticity of substitution. From the Euler equation (next slides):

$$E_t \Delta c_{t+1} = \frac{1}{\sigma}(r_t - \rho)$$

- Large σ: strong desire to smooth consumption; spending reacts little to the real rate.

φ: willingness to vary hours

- $1/\varphi$ is the Frisch elasticity of labor supply. From the labor FOC:

$$n_t = \frac{1}{\varphi}(w_t - \sigma c_t)$$

- Large φ: hours react little to the real wage; firms must pay much higher wages to get more hours.

Typical values: σ between 1 and 2 in macro models; Frisch elasticity around 0.5 in micro studies, often higher in macro models (Chetty et al., 2011, AER P&P).

What σ and φ imply for monetary policy

We will see this formally later in the lecture. Here is the intuition.

Large $\sigma \Rightarrow$ weaker effect of interest rates on spending

If households care a lot about smooth consumption, a rate increase changes their spending little. In the IS curve, $1/\sigma$ will be the response of demand to the real rate.

Large $\sigma + \varphi \Rightarrow$ inflation more sensitive to output

If workers need much higher wages to work more, firms' costs rise fast when output rises, and they raise prices more. In the Phillips curve, $\sigma + \varphi$ will enter the slope.

Solving the problem: the Lagrangian

We maximize utility subject to one budget constraint per period. Attach a multiplier λ_t to the period-t constraint:

$$\mathcal{L} = \mathrm{E}_0 \sum_{t=0}^{\infty} \beta^t \left\{ \frac{\mathrm{C}_t^{1-\sigma}}{1-\sigma} - \chi \frac{\mathrm{N}_t^{1+\varphi}}{1+\varphi} + \lambda_t \Big[(1+\mathrm{i}_{t-1})\mathrm{B}_{t-1} + \mathrm{W}_t \mathrm{N}_t + \mathrm{D}_t - \mathrm{P}_t \mathrm{C}_t - \mathrm{B}_t \Big] \right\}$$

- λ_t is the marginal utility of one extra euro in period t.
- Take the derivative with respect to each choice and set it to zero.
- Careful with B_t: it appears in period t (as savings, $-\mathrm{B}_t$) and in period $t+1$ (as income, $(1+\mathrm{i}_t)\mathrm{B}_t$, discounted by β).

First-order conditions (FOCs)

$$\mathrm{C}_t: \ \mathrm{C}_t^{-\sigma} = \lambda_t \mathrm{P}_t \qquad \mathrm{N}_t: \ \chi \mathrm{N}_t^{\varphi} = \lambda_t \mathrm{W}_t \qquad \mathrm{B}_t: \ \lambda_t = \beta(1+\mathrm{i}_t)\,\mathrm{E}_t \lambda_{t+1}$$

Step by step: look only at periods t and t + 1

The sum has infinitely many terms, but the choices made in period t (C_t, N_t, B_t) appear in only two of them: period t and period t + 1. We write those two out and ignore the rest.

$$\mathcal{L} = \ldots + \beta^t \Big\{ \frac{C_t^{1-\sigma}}{1-\sigma} - \chi \frac{N_t^{1+\varphi}}{1+\varphi} + \lambda_t \big[(1+i_{t-1}) B_{t-1} + W_t N_t + D_t - P_t C_t - B_t \big] \Big\}$$

$$+ \beta^{t+1} E_t \Big\{ \frac{C_{t+1}^{1-\sigma}}{1-\sigma} - \chi \frac{N_{t+1}^{1+\varphi}}{1+\varphi} + \lambda_{t+1} \big[(1+i_t) B_t + W_{t+1} N_{t+1} + D_{t+1} - P_{t+1} C_{t+1} - B_{t+1} \big] \Big\}$$

$$+ \ldots$$

- The "…" are all other periods: they contain no C_t, N_t or B_t, so their derivatives are zero.
- C_t and N_t appear only in period t.
- B_t appears in both: $-B_t$ today (money put aside) and $(1 + i_t)B_t$ tomorrow (money returned with interest).
- We stand in period t, so expectations of the future are E_t.

FOCs for consumption and labor

Consumption C_t. Only the period-t term contains C_t:

$$\frac{\partial \mathcal{L}}{\partial C_t} = \beta^t \Big[\underbrace{C_t^{-\sigma}}_{\text{from } \frac{C_t^{1-\sigma}}{1-\sigma}} \quad \underbrace{-\lambda_t P_t}_{\text{from } -\lambda_t P_t C_t} \Big] = 0 \quad \Longrightarrow \quad \boxed{C_t^{-\sigma} = \lambda_t P_t}$$

Labor N_t. Again only the period-t term:

$$\frac{\partial \mathcal{L}}{\partial N_t} = \beta^t \Big[\underbrace{-\chi N_t^{\varphi}}_{\text{from } -\chi \frac{N_t^{1+\varphi}}{1+\varphi}} \quad \underbrace{+\lambda_t W_t}_{\text{from } \lambda_t W_t N_t} \Big] = 0 \quad \Longrightarrow \quad \boxed{\chi N_t^{\varphi} = \lambda_t W_t}$$

- $\beta^t > 0$ multiplies everything, so it cancels.
- Consumption: the utility of one more unit of goods = the utility of the P_t euros it costs.
- Labor: the disutility of one more hour = the utility of the W_t euros it pays.

FOC for bonds: the only intertemporal choice

Bonds B_t. B_t appears in period t and in period $t+1$, so we get two pieces:

$$\frac{\partial \mathcal{L}}{\partial B_t} = \underbrace{\beta^t(-\lambda_t)}_{\text{period t: } -\lambda_t B_t} + \underbrace{\beta^{t+1} E_t\left[\lambda_{t+1}(1+i_t)\right]}_{\text{period t+1: } \lambda_{t+1}(1+i_t)B_t} = 0$$

Divide by β^t and move λ_t to the left. Since i_t is known in period t, it comes out of E_t:

$$\boxed{\lambda_t = \beta(1+i_t)\,E_t\lambda_{t+1}}$$

- Left: the utility lost by saving one euro today.
- Right: the utility gained tomorrow from the $1+i_t$ euros it returns, discounted by β.
- At the optimum, the household is indifferent between spending the last euro today and saving it.

The three FOCs together

First-order conditions

$$C_t:\ C_t^{-\sigma} = \lambda_t P_t \qquad N_t:\ \chi N_t^{\varphi} = \lambda_t W_t \qquad B_t:\ \lambda_t = \beta(1+i_t)\,E_t\lambda_{t+1}$$

- The same steps hold for every t, so these conditions hold in every period.
- The recipe: (i) find the terms that contain the variable, (ii) differentiate only those, (iii) set the result to zero.
- λ_t is not observed. Next, we use the FOCs to eliminate it:
  - C_t and B_t $\Rightarrow$ the Euler equation (saving vs. consumption over time);
  - C_t and N_t $\Rightarrow$ labor supply (work vs. consumption today).

Interpreting the first-order conditions

- λ_t converts euros into utility: one extra euro in period t raises utility by λ_t.
- C_t: the utility from one more unit of consumption, $C_t^{-\sigma}$, equals its cost in utility: P_t euros $\times$ λ_t per euro.
- N_t: the utility cost of one more hour, χN_t^{φ}, equals its benefit in utility: W_t euros $\times$ λ_t per euro.
- B_t: saving one euro costs λ_t in utility today; it returns $(1 + i_t)$ euros tomorrow, worth $\beta\lambda_{t+1}$ each in today's utility.

Next step

λ_t is not observed. We combine the FOCs to eliminate it.

Combining the first-order conditions

Consumption–leisure. Divide the N_t FOC by the C_t FOC; λ_t cancels:

Consumption–leisure (labor supply)

$$\frac{\chi N_t^{\varphi}}{C_t^{-\sigma}} = \frac{W_t}{P_t} \qquad \text{(marginal rate of substitution = real wage)}$$

Consumption–savings. From the C_t FOC, $\lambda_t = C_t^{-\sigma}/P_t$, and similarly for $t+1$. Substitute into the B_t FOC:

$$\frac{C_t^{-\sigma}}{P_t} = \beta(1+i_t)\,E_t\left[\frac{C_{t+1}^{-\sigma}}{P_{t+1}}\right]$$

Multiply both sides by P_t:

Consumption–savings (Euler equation)

$$C_t^{-\sigma} = \beta\,(1+i_t)\,E_t\left[C_{t+1}^{-\sigma}\,\frac{P_t}{P_{t+1}}\right]$$

The Euler equation: intuition

$$\underbrace{C_t^{-\sigma}}_{\text{utility lost by saving 1 unit today}} = \underbrace{\beta\,(1+i_t)\frac{P_t}{P_{t+1}}\,C_{t+1}^{-\sigma}}_{\text{utility gained tomorrow}}$$

- Give up one unit of consumption today. Save its value, P_t euros, in the bond.
- Tomorrow, receive $(1+i_t)P_t$ euros, which buy $(1+i_t)P_t/P_{t+1}$ units of goods: the gross real return.
- At the optimum, the household is indifferent: the two sides are equal.

What matters is the real rate

- The nominal rate appears only in this consumption–savings margin.
- Households care about $(1+i_t)P_t/P_{t+1}$: a higher nominal rate raises saving only if it is not offset by higher expected inflation.
- This is why sticky prices are needed for monetary policy to affect spending.

Why log-linearize?

- The optimality conditions are nonlinear: powers ($C_t^{-\sigma}$), ratios (P_t/P_{t+1}), expectations of nonlinear functions. In general they cannot be solved by hand.
- But business cycles are small fluctuations around a steady state: output, prices and rates move by a few percent.
- Close to that point, a nonlinear function is well approximated by a straight line — its tangent.

What we gain

- Linear equations that are easy to solve.
- Coefficients that are elasticities, with a direct economic meaning ($1/\sigma$, $1/\varphi$).
- Equations that can be estimated by regression — like the Taylor rule.

The cost: the approximation is accurate only near the steady state. Large shocks, or the zero lower bound, need nonlinear methods.

How we log-linearize

Step 1: take logs (exact)

Logs turn products into sums and powers into multiples:

$$\ln\left(X^a Y^b\right) = a \ln X + b \ln Y$$

e.g. $\ln C_t^{-\sigma} = -\sigma c_t$, with $c_t \equiv \ln C_t$.

Step 2: first-order Taylor approximation

Around $z = 0$: $\ln(1+z) \approx \ln 1 + \frac{1}{1+0} z = z$

z	$\ln(1+z)$	error
0.01	0.00995	0.00005
0.05	0.0488	0.0012

- The error is of order z^2 (second order); a first-order approximation ignores it. Quarterly rates are around 1%, so the error is tiny.
- Changes in logs are percentage changes: $\Delta c_{t+1} = 0.01$ means consumption grows by 1%.
- Step 3: deviations from steady state. Subtract the steady-state version of the equation; constants drop out.

Log-linearizing the Euler equation (1/2)

Start: $C_t^{-\sigma} = \beta\,(1+i_t)\,E_t\left[C_{t+1}^{-\sigma}\,\frac{P_t}{P_{t+1}}\right]$

1. Take logs of both sides; the log of a product is the sum of the logs (exact):

$$-\sigma\, c_t = \ln\beta + \ln(1+i_t) + \ln E_t\left[C_{t+1}^{-\sigma}\,\frac{P_t}{P_{t+1}}\right]$$

2. Linearize term by term (a): the expectation term. Move E_t inside the log:

$$\ln E_t[X] \approx E_t[\ln X]$$

The difference depends on the variance of X and is second order (we ignore uncertainty to first order). So:

$$-\sigma\, c_t = \ln\beta + \ln(1+i_t) - \sigma\, E_t c_{t+1} - E_t \ln\frac{P_{t+1}}{P_t}$$

Log-linearizing the Euler equation (2/2)

3. Linearize term by term (b): the remaining terms.
   - $-\sigma c_t$ and $-\sigma E_t c_{t+1}$ are already linear.
   - $\ln \beta = -\rho$ is a constant. $\rho \equiv -\ln \beta$ is the discount rate ($\approx 1\%$ per quarter when $\beta = 0.99$).
   - The rates use $\ln(1 + z) \approx z$: $\ln(1 + i_t) \approx i_t$ and $\ln(P_{t+1}/P_t) = \ln(1 + \pi_{t+1}) \approx \pi_{t+1}$.

$$-\sigma\, c_t = -\rho + i_t - E_t \pi_{t+1} - \sigma\, E_t c_{t+1}$$

4. Divide by $-\sigma$ and rearrange.

Linearized Euler equation

$$c_t = E_t c_{t+1} - \frac{1}{\sigma}\left(i_t - E_t \pi_{t+1} - \rho\right)$$

Steps 2 and 3 are both first-order approximations; their order does not matter.

The Euler equation in steady state: the natural real rate

Steady state: Consumption is constant, $c_t = E_t c_{t+1} = \bar{c}$; inflation and the interest rate are constant, $\bar{\pi}$ and $\bar{i}$. The Euler equation becomes

$$\bar{c} = \bar{c} - \frac{1}{\sigma}\left(\bar{i} - \bar{\pi} - \rho\right) \quad \Longrightarrow \quad \bar{r} \equiv \bar{i} - \bar{\pi} = \rho$$

In steady state, the real interest rate equals the discount rate.

- Why? ρ measures impatience. A very impatient household (high ρ) needs a high real interest rate to give up consumption today and save. If $r > \rho$, saving pays more than impatience costs, so the household saves and consumption grows. If $r < \rho$, it borrows and consumption falls. Consumption stays constant only when $r = \rho$.
- This rate is the natural (or neutral) real interest rate, often written r^*: the real rate consistent with a stable economy, neither stimulating nor restraining demand.
- It is determined by preferences (β), not by the central bank.

The Euler equation in deviations from steady state

Subtract the steady-state equation from the linearized one:

$$(c_t - \bar{c}) = E_t(c_{t+1} - \bar{c}) - \tfrac{1}{\sigma}\Big[(i_t - E_t\pi_{t+1}) - \rho\Big]$$

- $\bar{c}$ appears on both sides and cancels: the equation has the same form in deviations.
- $(i_t - E_t\pi_{t+1}) - \rho$ is the deviation of the real rate from its natural level.

From now on, c_t denotes the percentage deviation of consumption from steady state

$$c_t = E_t c_{t+1} - \frac{1}{\sigma}(r_t - \rho), \qquad r_t \equiv i_t - E_t\pi_{t+1} \text{ (expected real interest rate)}$$

- Link to the Taylor rule: $i_t = r^* + \pi^* + \phi_\pi(\pi_t - \pi^*) + \phi_x x_t$. With inflation at target and a zero gap, $i_t = r^* + \pi^*$, so the real rate equals r^*: policy is neutral. In our model $\pi^* = 0$ and $r^* = \rho$, so the constant is ρ. (Taylor, 1993: $r^* = 2\%$.)
- Later, the natural rate will also move with shocks (r_t^n); ρ is its steady-state value.

Interpreting the Euler equation (1/2): ρ and the steady state

- $\rho \equiv -\ln \beta \approx 1 - \beta$: the rate of time preference, i.e. impatience: how much less the household values utility one period ahead.
- With $\beta = 0.99$, $\rho \approx 1\%$ per quarter, about 4% per year.
- Steady state: $r = \rho$. To give up consumption today, the household must earn a real return that compensates its impatience.
- A more impatient household (lower β, higher ρ) requires a higher real rate to save.

Away from steady state

- $r_t > \rho$ more than compensates impatience: saving is attractive; consumption today falls relative to tomorrow.
- $r_t < \rho$: borrowing is attractive; consumption today rises.
- $1/\sigma$ measures how strongly consumption responds.

Interpreting the Euler equation (2/2): expectations

$$c_t = E_t c_{t+1} - \frac{1}{\sigma}(r_t - \rho)$$

- Expectations matter: higher expected future consumption raises consumption today (households want to smooth).
- Iterating forward: substitute for $E_t c_{t+1}$ using the same equation one period ahead, then for $E_t c_{t+2}$, and so on:

$$c_t = -\tfrac{1}{\sigma} \sum_{k \geq 0} E_t(r_{t+k} - \rho) + \lim_{k \to \infty} E_t c_{t+k}$$

- Consumption today depends on the whole path of expected real rates, not only today's rate.
- This is why central banks care about expected future policy, and why communication (forward guidance) can matter.

This is the seed of the IS curve

Once we impose that consumption equals output, this becomes the demand side of the model.

Labor supply, in logs and deviations

Start from the consumption–leisure condition: $\chi N_t^{\varphi} C_t^{\sigma} = W_t/P_t$

Take logs. No approximation is needed — only products and powers. With $w_t \equiv \ln(W_t/P_t)$:

$$w_t = \ln \chi + \sigma\, c_t + \varphi\, n_t$$

Deviations from steady state. The same equation holds in steady state: $\bar{w} = \ln \chi + \sigma \bar{c} + \varphi \bar{n}$. Subtract it:

$$(w_t - \bar{w}) = \sigma\,(c_t - \bar{c}) + \varphi\,(n_t - \bar{n})$$

The constant $\ln \chi$ drops out and the coefficients are unchanged. As with consumption, from now on w_t and n_t denote percentage deviations from steady state.

Labor supply

$$w_t = \sigma\, c_t + \varphi\, n_t \qquad \Longleftrightarrow \qquad n_t = \tfrac{1}{\varphi}\,(w_t - \sigma\, c_t)$$

Interpreting labor supply

$$w_t = \sigma\, c_t + \varphi\, n_t$$

- To work more hours (n_t up), households require a higher real wage; how much higher depends on φ.
- Smaller φ: labor supply is more elastic (the Frisch elasticity $1/\varphi$ is larger). A small wage increase brings many more hours. As $\varphi \to 0$, households supply any hours at the going wage.
- When households consume more (c_t up), they value leisure more and also require a higher wage (a wealth effect), governed by σ.
- Read it as the wage that firms must pay to get n_t hours, given consumption c_t.

Where this is going

Firms' costs depend on wages. With this equation, wages — and hence firms' costs — will move with output: the smaller φ, the more slowly. This is the link from the real economy to inflation.

Households: summary

Euler equation: $c_t = E_t c_{t+1} - \frac{1}{\sigma}\left(i_t - E_t \pi_{t+1} - \rho\right)$ Labor supply: $w_t = \sigma c_t + \varphi n_t$

Firms

Why two layers of firms?

- For firms to set prices, they need market power: each must sell a good that is not a perfect substitute for the others.
- But households buy one final good. The two layers make this tractable:
  - Wholesale firms, indexed by $i \in [0, 1]$, each produce a different variety (e.g. brands) and set its price.
  - Retail firms buy all varieties and combine them into the final good, in perfect competition.
- The retail firms' problem gives the demand curve each wholesale firm faces.
- The wholesale firms' problem gives pricing: first with flexible prices, then with costs of changing prices.

Notation

i runs over all varieties (inside integrals); j denotes one particular variety or firm, whose decision we study.

Retail firms: the CES aggregator

A retail firm combines the varieties with a constant elasticity of substitution (CES, Dixit–Stiglitz) technology:

$$Y_t = \left[\int_0^1 Y_{it}^{\frac{\varepsilon-1}{\varepsilon}} \, di \right]^{\frac{\varepsilon}{\varepsilon-1}}, \qquad \varepsilon > 1$$

- Y_{it}: quantity of variety i used as an input; Y_t: output of the final good.
- The integral is a "sum" over a continuum of varieties.
- ε: how easily the varieties can be substituted for one another.
  - Large ε: varieties are close substitutes; each firm has little market power.
  - $\varepsilon \to \infty$: perfect substitutes, perfect competition.
- If all varieties are used in the same quantity, $Y_{it} = Y_t$: the aggregator just adds them up.

The retail firm's problem

- The retail firm is perfectly competitive: it takes the price of the final good P_t and the prices of all varieties P_{it} as given.
- It chooses how much of each variety to buy, to maximize profits:

$$\max_{\{Y_{it}\}_{i\in[0,1]}} \underbrace{P_t Y_t}_{\text{revenue}} - \underbrace{\int_0^1 P_{it} Y_{it}\, di}_{\text{cost of inputs}}, \qquad Y_t = \left[\int_0^1 Y_{it}^{\frac{\varepsilon-1}{\varepsilon}} di\right]^{\frac{\varepsilon}{\varepsilon-1}}$$

What we will get

- Take the FOC for one variety, j: the demand for good j, as a function of its price relative to the others. The same holds for every j.
- The price index P_t: the price of the final good.

Deriving the demand for good j (1/2)

Write profits with Y_t substituted in, so that we see every place where Y_{jt} appears:

$$\Pi_t = \underbrace{P_t \left[\int_0^1 Y_{it}^{\frac{\varepsilon-1}{\varepsilon}} \, di\right]^{\frac{\varepsilon}{\varepsilon-1}}}_{\text{revenue } P_t Y_t} - \underbrace{\int_0^1 P_{it} Y_{it} \, di}_{\text{cost of inputs}}$$

- Think of each integral as a sum over varieties: only the term $i = j$ contains Y_{jt}.
- Revenue: chain rule — derivative of the outer bracket $\times$ derivative of the term $i = j$ inside.
- Cost: the term $i = j$ is $P_{jt} Y_{jt}$, with derivative P_{jt}.

Differentiate with respect to Y_{jt} and set to zero:

$$\frac{\partial \Pi_t}{\partial Y_{jt}} = P_t \underbrace{\frac{\varepsilon}{\varepsilon-1}\left[\int_0^1 Y_{it}^{\frac{\varepsilon-1}{\varepsilon}} \, di\right]^{\frac{\varepsilon}{\varepsilon-1}-1}}_{\text{outer: } \frac{d}{dz} z^{\frac{\varepsilon}{\varepsilon-1}}} \underbrace{\frac{\varepsilon-1}{\varepsilon} Y_{jt}^{\frac{\varepsilon-1}{\varepsilon}-1}}_{\text{inner: } \frac{d}{dY_{jt}} Y_{jt}^{\frac{\varepsilon-1}{\varepsilon}}} - \underbrace{P_{jt}}_{\text{cost}} = 0$$

Deriving the demand for good j (2/2)

Simplify. The constants cancel, $\frac{\varepsilon}{\varepsilon-1}\cdot\frac{\varepsilon-1}{\varepsilon}=1$, and the exponents are $\frac{\varepsilon}{\varepsilon-1}-1=\frac{1}{\varepsilon-1}$ and $\frac{\varepsilon-1}{\varepsilon}-1=-\frac{1}{\varepsilon}$:

$$\mathrm{P_t}\left[\int_0^1 \mathrm{Y_{it}}^{\frac{\varepsilon-1}{\varepsilon}}\,\mathrm{di}\right]^{\frac{1}{\varepsilon-1}}\mathrm{Y_{jt}}^{-\frac{1}{\varepsilon}}=\mathrm{P_{jt}}$$

Recognize $\mathrm{Y_t}$. Write the exponent as $\frac{1}{\varepsilon-1}=\frac{\varepsilon}{\varepsilon-1}\cdot\frac{1}{\varepsilon}$:

$$\mathrm{P_t}\Bigg(\underbrace{\left[\int_0^1 \mathrm{Y_{it}}^{\frac{\varepsilon-1}{\varepsilon}}\,\mathrm{di}\right]^{\frac{\varepsilon}{\varepsilon-1}}}_{=\mathrm{Y_t}}\Bigg)^{\frac{1}{\varepsilon}}\mathrm{Y_{jt}}^{-\frac{1}{\varepsilon}}=\mathrm{P_{jt}}\quad\Longrightarrow\quad\mathrm{P_t}\,\mathrm{Y_t}^{\frac{1}{\varepsilon}}\,\mathrm{Y_{jt}}^{-\frac{1}{\varepsilon}}=\mathrm{P_{jt}}$$

Solve for $\mathrm{Y_{jt}}$. Divide by $\mathrm{P_t}\mathrm{Y_t}^{1/\varepsilon}$, then raise both sides to the power $-\varepsilon$:

$$\mathrm{Y_{jt}}^{-\frac{1}{\varepsilon}}=\frac{\mathrm{P_{jt}}}{\mathrm{P_t}}\,\mathrm{Y_t}^{-\frac{1}{\varepsilon}}\quad\Longrightarrow\quad\boxed{\mathrm{Y_{jt}}=\left(\frac{\mathrm{P_{jt}}}{\mathrm{P_t}}\right)^{-\varepsilon}\mathrm{Y_t}}$$

Interpreting demand, and the price index

$$Y_{jt} = \left(\frac{P_{jt}}{P_t}\right)^{-\varepsilon} Y_t$$

- Demand for good j falls when its price rises relative to the average price P_t.
- ε is the price elasticity of demand: a 1% higher relative price lowers demand by ε%.
- Demand rises one-for-one with total output Y_t.

The price index. Retail is competitive, so profits are zero: $P_t Y_t = \int_0^1 P_{it} Y_{it}\, di$. Substitute the demand for each variety i:

$$P_t Y_t = \int_0^1 P_{it} \left(\frac{P_{it}}{P_t}\right)^{-\varepsilon} Y_t\, di \quad \Longrightarrow \quad P_t^{1-\varepsilon} = \int_0^1 P_{it}^{1-\varepsilon}\, di$$

Aggregate price index

$$P_t = \left[\int_0^1 P_{it}^{1-\varepsilon}\, di\right]^{\frac{1}{1-\varepsilon}}$$

If all firms charge the same price: $P_t = P_{it}$.

Wholesale firms: technology and marginal cost

- Firm j produces with labor only: $Y_{jt} = A_t N_{jt}$, where A_t is productivity (common to all firms).
- To produce one more unit, the firm needs $1/A_t$ more hours, each costing the nominal wage W_t.
- Nominal marginal cost: W_t/A_t. Divide by P_t to express it in units of goods:

Real marginal cost

$$mc_t = \frac{W_t/P_t}{A_t} \qquad \text{in logs:} \quad \ln mc_t = w_t - a_t$$

- The same for all firms: they hire in the same labor market and use the same technology.
- Marginal cost rises when real wages rise and falls when productivity rises.
- Marginal cost will be the key link between the real economy and inflation.

Monopolistic competition

Each wholesale firm j is a monopolist for its own variety, but one of many firms:

- It sets its price P_{jt} and faces a downward-sloping demand curve, $Y_{jt} = (P_{jt}/P_t)^{-\varepsilon} Y_t$.
- It is small: its own price does not affect the price index P_t or aggregate output Y_t, so it takes them as given.
- No strategic interaction: the firm does not need to guess how other firms will react.

	Perfect competition	Monopolistic competition	Monopoly
Number of firms	many	many	one
Product	identical	differentiated	unique
Sets its price?	no	yes	yes
Affects aggregates?	no	no	yes
Price	= marginal cost	markup × marginal cost	markup × marginal cost

Dixit and Stiglitz (1977), AER; course notes, Chapter 22.

Why the NK model needs monopolistic competition

- Sticky prices require price setters. Under perfect competition, firms take the market price as given; there is no price for them to keep fixed.

- Price above marginal cost. With a markup, a firm is willing to sell more at its current price when demand rises: each extra unit is still profitable.

- Output is demand-determined in the short run. When prices are sticky, firms meet the demand they face at their posted price. So a change in demand — for example, from a change in the real interest rate — changes output.

Next

How does firm j choose its price? First with flexible prices (no cost of changing it), then with sticky prices.

Benchmark: flexible prices and the markup

If changing prices is free, firm j maximizes profits taking its demand curve into account:

$$\max_{P_{jt}} \underbrace{P_{jt}\, Y_{jt}}_{\text{revenue}} - \underbrace{\frac{W_t}{A_t}\, Y_{jt}}_{\text{cost}} \qquad \text{subject to} \qquad Y_{jt} = \left(\frac{P_{jt}}{P_t}\right)^{-\varepsilon} Y_t$$

Exercise

Substitute the demand curve into profits, differentiate with respect to P_{jt}, and show that $P_{jt} = \mathcal{M}\, \frac{W_t}{A_t}$, with $\mathcal{M} \equiv \frac{\varepsilon}{\varepsilon - 1}$, or in real terms $\frac{P_{jt}}{P_t} = \mathcal{M}\, mc_t$.

- The price is a markup $\mathcal{M} > 1$ over marginal cost: to sell one more unit, the firm must lower its price on all units, so it sells less and charges more than under perfect competition.
- $\mathcal{M}$ depends only on ε: close substitutes mean a small markup; as $\varepsilon \to \infty$, $\mathcal{M} \to 1$. With $\varepsilon = 6$, prices are 20% above marginal cost.
- In equilibrium $P_{jt} = P_t$, so $mc_t = 1/\mathcal{M}$: with flexible prices, real marginal cost is constant.

Prices are sticky in the data

- Micro data on individual prices: prices change infrequently.
  - Euro area: on average about 15% of prices change in a given month — a typical price lasts about a year.
  - United States: prices change more often, about 25% per month.
- Prices of services and processed food change less often; energy and fresh food more often.

Why do firms not change prices continuously?

- Menu costs: any cost incurred in the act of changing a price — printing menus, relabeling, managerial time, upsetting customers.
- Fixed menu costs: independent of the size of the change.
- Convex costs: larger changes cost disproportionately more.

Dhyne et al. (2006), "Price Changes in the Euro Area and the United States," Journal of Economic Perspectives; Nakamura and Steinsson (2008), QJE.

Rotemberg (1982): quadratic adjustment cost

Cost of changing the price, paid by firm j in units of goods

$$\frac{\psi}{2}\left(\frac{P_{jt}}{P_{j,t-1}}-1\right)^2 Y_t$$

- $\frac{P_{jt}}{P_{j,t-1}}-1$: the percentage change in the firm's price.
- $\psi > 0$ measures the importance of the cost; $\psi = 0$ recovers flexible prices.
- No change in price, no cost. Increases and decreases cost the same.
- The cost is real: it uses up goods, scaled by the size of the economy Y_t.
- Convexity gives an incentive to adjust gradually: two changes of 5% cost $2 \times \frac{\psi}{2}(0.05)^2$, half the cost of one change of 10%, $\frac{\psi}{2}(0.10)^2$. Prices become sticky.

Rotemberg, J. (1982), "Sticky Prices in the United States," JPE. The alternative, Calvo (1983) pricing — each period only a random fraction of firms can change price — gives the same linear Phillips curve.

The firm's problem becomes dynamic

- A price chosen today affects the adjustment cost tomorrow, because tomorrow's change is measured from today's price.
- The firm can no longer choose each period's price separately: it must look ahead.
- It maximizes the present value of all current and future real profits.

Discounting future profits

- Profits in period t are multiplied by a discount factor Q_t: one unit of goods in period t is worth Q_t units today. $Q_0 = 1$.
- Q_t plays the same role as β^t for households: future profits count less than profits today.
- What determines Q_t? We will come back to it after we derive the Phillips curve.

The wholesale firm's dynamic problem

$$\max_{\{P_{jt}\}} E_0 \sum_{t=0}^{\infty} Q_t \left[\underbrace{\left(\frac{P_{jt}}{P_t} - mc_t \right) \left(\frac{P_{jt}}{P_t} \right)^{-\varepsilon} Y_t}_{\text{profits from sales}} - \underbrace{\frac{\psi}{2} \left(\frac{P_{jt}}{P_{j,t-1}} - 1 \right)^2 Y_t}_{\text{cost of changing price}} \right]$$

- First term: the same as with flexible prices — profit per unit times units sold.
- Second term: the adjustment cost.
- The firm takes P_t, mc_t, Y_t and Q_t as given; it chooses only its own price.

A useful rewrite. Multiply out the profits from sales: $\left(\frac{P_{jt}}{P_t} \right) \left(\frac{P_{jt}}{P_t} \right)^{-\varepsilon} = \left(\frac{P_{jt}}{P_t} \right)^{1-\varepsilon}$, so

$$\left(\frac{P_{jt}}{P_t} - mc_t \right) \left(\frac{P_{jt}}{P_t} \right)^{-\varepsilon} Y_t = \underbrace{\left(\frac{P_{jt}}{P_t} \right)^{1-\varepsilon} Y_t}_{\text{revenue}} - \underbrace{mc_t \left(\frac{P_{jt}}{P_t} \right)^{-\varepsilon} Y_t}_{\text{cost of production}}$$

Each piece is now a simple power of P_{jt}.

Step by step: look only at periods t and t + 1

As for households, P_{jt} appears in only two terms of the sum. We write those two out:

$$\Pi = \ldots + Q_t \Big[\underbrace{\left(\frac{P_{jt}}{P_t}\right)^{1-\varepsilon} Y_t}_{\text{(1) revenue}} - \underbrace{mc_t \left(\frac{P_{jt}}{P_t}\right)^{-\varepsilon} Y_t}_{\text{(2) production cost}} - \underbrace{\frac{\psi}{2}\left(\frac{P_{jt}}{P_{j,t-1}} - 1\right)^2 Y_t}_{\text{(3) adjustment cost today}} \Big]$$

$$+ E_t\, Q_{t+1} \Big[\underbrace{\left(\frac{P_{j,t+1}}{P_{t+1}}\right)^{1-\varepsilon} Y_{t+1} - mc_{t+1} \left(\frac{P_{j,t+1}}{P_{t+1}}\right)^{-\varepsilon} Y_{t+1}}_{\text{no } P_{jt}} - \underbrace{\frac{\psi}{2}\left(\frac{P_{j,t+1}}{P_{jt}} - 1\right)^2 Y_{t+1}}_{\text{(4) adjustment cost tomorrow}} \Big] + \ldots$$

- Period t: P_{jt} appears in three terms — revenue (1), production cost (2), and today's adjustment cost (3).
- Period $t + 1$: P_{jt} appears only in tomorrow's adjustment cost (4), because tomorrow's price change is measured from today's price. Sales in $t + 1$ depend on $P_{j,t+1}$, not on P_{jt}.
- All other periods contain no P_{jt}: their derivatives are zero.

So the FOC has four pieces. We take them one at a time.

Pieces (1) and (2): profits from sales in period t

$$\Pi = \ldots + Q_t \left[\left(\frac{P_{jt}}{P_t} \right)^{1-\varepsilon} Y_t - mc_t \left(\frac{P_{jt}}{P_t} \right)^{-\varepsilon} Y_t - \frac{\psi}{2} \left(\frac{P_{jt}}{P_{j,t-1}} - 1 \right)^2 Y_t \right]$$
$$+ E_t \, Q_{t+1} \left[\left(\frac{P_{j,t+1}}{P_{t+1}} \right)^{1-\varepsilon} Y_{t+1} - mc_{t+1} \left(\frac{P_{j,t+1}}{P_{t+1}} \right)^{-\varepsilon} Y_{t+1} - \frac{\psi}{2} \left(\frac{P_{j,t+1}}{P_{jt}} - 1 \right)^2 Y_{t+1} \right] + \ldots$$

Rule: $\frac{d}{dP_{jt}} \left(\frac{P_{jt}}{P_t} \right)^k = k \left(\frac{P_{jt}}{P_t} \right)^{k-1} \cdot \frac{1}{P_t}$ (power rule, times the derivative of the inside, $\frac{1}{P_t}$).

(1) $\frac{\partial}{\partial P_{jt}} \left[\left(\frac{P_{jt}}{P_t} \right)^{1-\varepsilon} Y_t \right] = (1-\varepsilon) \left(\frac{P_{jt}}{P_t} \right)^{-\varepsilon} \frac{Y_t}{P_t}$

(2) $\frac{\partial}{\partial P_{jt}} \left[- mc_t \left(\frac{P_{jt}}{P_t} \right)^{-\varepsilon} Y_t \right] = \varepsilon \, mc_t \left(\frac{P_{jt}}{P_t} \right)^{-\varepsilon-1} \frac{Y_t}{P_t}$

- (1) is negative ($\varepsilon > 1$): with elastic demand, a higher price lowers revenue, because the lost sales outweigh the higher price per unit.
- (2) is positive: a higher price means fewer units sold, so lower production costs. (The minus sign in front times the exponent $-\varepsilon$ gives $+\varepsilon$.)
- With flexible prices ($\psi = 0$), (1) + (2) = 0 is the whole FOC: it gives the markup rule $\frac{P_{jt}}{P_t} = \frac{\varepsilon}{\varepsilon - 1} mc_t$.

Pieces (3) and (4): the adjustment costs, today and tomorrow

$$\Pi = \ldots + Q_t\left[\left(\frac{P_{jt}}{P_t}\right)^{1-\varepsilon} Y_t - mc_t\left(\frac{P_{jt}}{P_t}\right)^{-\varepsilon} Y_t - \frac{\psi}{2}\left(\frac{P_{jt}}{P_{j,t-1}} - 1\right)^2 Y_t\right]$$
$$+ E_t\, Q_{t+1}\left[\left(\frac{P_{j,t+1}}{P_{t+1}}\right)^{1-\varepsilon} Y_{t+1} - mc_{t+1}\left(\frac{P_{j,t+1}}{P_{t+1}}\right)^{-\varepsilon} Y_{t+1} - \frac{\psi}{2}\left(\frac{P_{j,t+1}}{P_{jt}} - 1\right)^2 Y_{t+1}\right] + \ldots$$

Rule: $\frac{d}{dz}\frac{\psi}{2}(z-1)^2 = \psi(z-1)$, times the derivative of the inside z.

(3) $$\frac{\partial}{\partial P_{jt}}\left[-\frac{\psi}{2}\left(\frac{P_{jt}}{P_{j,t-1}} - 1\right)^2 Y_t\right] = -\psi\left(\frac{P_{jt}}{P_{j,t-1}} - 1\right)\frac{Y_t}{P_{j,t-1}}$$ inside: $\frac{d}{dP_{jt}}\frac{P_{jt}}{P_{j,t-1}} = \frac{1}{P_{j,t-1}}$

(4) $$\frac{\partial}{\partial P_{jt}}\left[-\frac{\psi}{2}\left(\frac{P_{j,t+1}}{P_{jt}} - 1\right)^2 Y_{t+1}\right] = -\psi\left(\frac{P_{j,t+1}}{P_{jt}} - 1\right)\left(-\frac{P_{j,t+1}}{P_{jt}^2}\right) Y_{t+1}$$ inside: $\frac{d}{dP_{jt}}\frac{P_{j,t+1}}{P_{jt}} = -\frac{P_{j,t+1}}{P_{jt}^2}$

- (3): if the firm is raising its price, a higher P_{jt} means a bigger change today, and a bigger cost.
- (4): P_{jt} is in the denominator; the two minus signs make the term positive. A higher P_{jt} means a smaller change is needed tomorrow, so tomorrow's cost falls.

Putting the pieces together

Add the four pieces, each with its weight (Q_t for period t, $E_t Q_{t+1}$ for period $t+1$), set the sum to zero, and divide by Q_t:

$$\underbrace{(1-\varepsilon)\left(\frac{P_{jt}}{P_t}\right)^{-\varepsilon}\frac{Y_t}{P_t}+\varepsilon\, mc_t\left(\frac{P_{jt}}{P_t}\right)^{-\varepsilon-1}\frac{Y_t}{P_t}}_{\text{(a) = (1) + (2): sales}}\underbrace{-\psi\left(\frac{P_{jt}}{P_{j,t-1}}-1\right)\frac{Y_t}{P_{j,t-1}}}_{\text{(b) = (3): cost today}}\underbrace{+E_t\left[\frac{Q_{t+1}}{Q_t}\,\psi\left(\frac{P_{j,t+1}}{P_{jt}}-1\right)\frac{P_{j,t+1}}{P_{jt}^2}\,Y_{t+1}\right]}_{\text{(c) = (4): cost tomorrow}}=0$$

- Dividing by Q_t leaves period-t terms unweighted; the period-$t+1$ term is weighted by Q_{t+1}/Q_t, the value of tomorrow's goods in today's goods.
- (a): the only term without ψ. On its own (flexible prices, $\psi = 0$), it gives the desired markup rule: next slide.
- (b): raising the price today increases today's adjustment cost.
- (c): raising the price today reduces the increase needed tomorrow, and so tomorrow's cost.

The firm trades off three things

Being close to its desired (flexible-price) markup today, the cost of changing price today, and the cost of changing price tomorrow. If it expects to raise prices tomorrow, it starts today.

Special case: flexible prices ($\psi = 0$)

If changing prices is free, $\psi = 0$: terms (b) and (c) disappear and only (a) remains:

$$(1-\varepsilon)\left(\frac{P_{jt}}{P_t}\right)^{-\varepsilon}\frac{Y_t}{P_t} + \varepsilon\, mc_t \left(\frac{P_{jt}}{P_t}\right)^{-\varepsilon-1}\frac{Y_t}{P_t} = 0$$

Divide by $\left(\frac{P_{jt}}{P_t}\right)^{-\varepsilon-1}\frac{Y_t}{P_t}$ and solve:

$$(1-\varepsilon)\frac{P_{jt}}{P_t} + \varepsilon\, mc_t = 0 \quad \Longrightarrow \quad \frac{P_{jt}}{P_t} = \frac{\varepsilon}{\varepsilon-1}\, mc_t \quad \Longleftrightarrow \quad P_{jt} = \frac{\varepsilon}{\varepsilon-1}\,\frac{W_t}{A_t}$$

(This is the answer to the flexible-price exercise.)

Markup: price over nominal marginal cost

$$\text{markup} \equiv \frac{P_{jt}}{W_t/A_t} \qquad \text{With flexible prices:} \quad \text{markup} = \mathcal{M} \equiv \frac{\varepsilon}{\varepsilon-1} > 1$$

The flexible-price markup: intuition

- What determines $\mathcal{M}$? Only ε. Close substitutes (high ε) mean little market power and a small markup; as $\varepsilon \to \infty$, $\mathcal{M} \to 1$.
- Example: with $\varepsilon = 6$, $\mathcal{M} = 6/5 = 1.2$: prices are 20% above marginal cost.
- Constant: $\mathcal{M}$ does not depend on demand, wages or productivity. When marginal cost rises by 5%, the firm raises its price by 5%, and the markup is unchanged.

The firm's target

$\mathcal{M}$ is the markup the firm wants to charge: its desired markup. With flexible prices it always gets it. With sticky prices ($\psi > 0$), it may not.

Sticky prices: actual and desired markups

- Desired markup: $\mathcal{M} = \frac{\varepsilon}{\varepsilon - 1}$, the flexible-price markup. It is constant.
- Actual markup: the same definition, price over nominal marginal cost, at the prices firms actually charge. If all firms charge the same price ($P_{jt} = P_t$, formalized shortly):
- Flexible prices: $\mathcal{M}_t = \mathcal{M}$ every period, so $mc_t = 1/\mathcal{M}$ is constant.
- Sticky prices: $\mathcal{M}_t$ can differ from $\mathcal{M}$, because adjusting the price is costly.
- Steady state: with zero inflation, $\bar{\mathcal{M}} = \mathcal{M}$. With positive trend inflation, firms never fully catch up with rising costs, so $\bar{\mathcal{M}} < \mathcal{M}$.

Why this matters for inflation

- Demand rises $\Rightarrow$ wages and marginal cost rise $\Rightarrow$ prices do not adjust fully $\Rightarrow$ the actual markup falls below $\mathcal{M}$.
- Firms raise prices gradually to restore their markup: this is inflation.
- Demand falls $\Rightarrow$ the opposite: the actual markup rises above $\mathcal{M}$. A higher markup is not better: $\mathcal{M}$ maximizes profits, and above it the firm loses more in sales than it gains per unit. So firms cut prices gradually.

Symmetric equilibrium

The firm's FOC, from before:

$$\underbrace{(1-\varepsilon)\left(\frac{P_{jt}}{P_t}\right)^{-\varepsilon}\frac{Y_t}{P_t}+\varepsilon\,mc_t\left(\frac{P_{jt}}{P_t}\right)^{-\varepsilon-1}\frac{Y_t}{P_t}}_{\text{(a)}}\ \underbrace{-\psi\left(\frac{P_{jt}}{P_{j,t-1}}-1\right)\frac{Y_t}{P_{j,t-1}}}_{\text{(b)}}\ \underbrace{+E_t\left[\frac{Q_{t+1}}{Q_t}\,\psi\left(\frac{P_{j,t+1}}{P_{jt}}-1\right)\frac{P_{j,t+1}}{P_{jt}^2}\,Y_{t+1}\right]}_{\text{(c)}}=0$$

- All firms face the same problem, so they choose the same price: $P_{jt}=P_t$ for all j, in every period.
- Define (net) inflation by $1+\pi_t\equiv P_t/P_{t-1}$, so that $P_t/P_{t-1}-1=\pi_t$.

Impose symmetry term by term:

- (a): $\left(\frac{P_{jt}}{P_t}\right)^{-\varepsilon}=\left(\frac{P_{jt}}{P_t}\right)^{-\varepsilon-1}=1$, so $(1-\varepsilon)\frac{Y_t}{P_t}+\varepsilon\,mc_t\frac{Y_t}{P_t}$.
- (b): $\frac{P_{jt}}{P_{j,t-1}}-1=\pi_t$, so $-\psi\,\pi_t\,\frac{Y_t}{P_{t-1}}$.
- (c): $\frac{P_{j,t+1}}{P_{jt}}-1=\pi_{t+1}$ and $\frac{P_{t+1}}{P_t^2}=\frac{1+\pi_{t+1}}{P_t}$, so $E_t\left[\frac{Q_{t+1}}{Q_t}\,\psi\,\pi_{t+1}(1+\pi_{t+1})\,\frac{Y_{t+1}}{P_t}\right]$.

The FOC with $P_{jt}=P_t$

$$(1-\varepsilon)\frac{Y_t}{P_t}+\varepsilon\,mc_t\frac{Y_t}{P_t}-\psi\,\pi_t\,\frac{Y_t}{P_{t-1}}+E_t\left[\frac{Q_{t+1}}{Q_t}\,\psi\,\pi_{t+1}(1+\pi_{t+1})\,\frac{Y_{t+1}}{P_t}\right]=0$$

The FOC in terms of inflation: the nonlinear Phillips curve

Multiply every term by P_t/Y_t, to remove the price level and express everything in inflation rates:

Term	With $P_{jt} = P_t$	Times P_t/Y_t
(a)	$(1-\varepsilon)\frac{Y_t}{P_t} + \varepsilon\, mc_t \frac{Y_t}{P_t}$	$(1-\varepsilon) + \varepsilon\, mc_t$
(b)	$-\psi\, \pi_t \frac{Y_t}{P_{t-1}}$	$-\psi\, \pi_t \frac{P_t}{P_{t-1}} = -\psi\, \pi_t(1+\pi_t)$
(c)	$E_t\left[\frac{Q_{t+1}}{Q_t}\psi\, \pi_{t+1}(1+\pi_{t+1})\frac{Y_{t+1}}{P_t}\right]$	$E_t\left[\frac{Q_{t+1}}{Q_t}\frac{Y_{t+1}}{Y_t}\psi\, \pi_{t+1}(1+\pi_{t+1})\right]$

Move (b) to the left-hand side:

The New Keynesian Phillips curve (nonlinear)

$$\psi\, \pi_t(1+\pi_t) = (1-\varepsilon) + \varepsilon\, mc_t + E_t\left[\frac{Q_{t+1}}{Q_t}\,\frac{Y_{t+1}}{Y_t}\,\psi\, \pi_{t+1}(1+\pi_{t+1})\right]$$

The Phillips curve in terms of markups

Rewrite the first term of the nonlinear Phillips curve, using $\mathcal{M} = \frac{\varepsilon}{\varepsilon - 1}$ and $\text{mc}_t = 1/\mathcal{M}_t$:

$$(1 - \varepsilon) + \varepsilon\, \text{mc}_t = \varepsilon \left(\text{mc}_t - \frac{\varepsilon - 1}{\varepsilon} \right) = \varepsilon \left(\frac{1}{\mathcal{M}_t} - \frac{1}{\mathcal{M}} \right)$$

The Phillips curve

$$\psi\, \pi_t (1 + \pi_t) = \varepsilon \left(\frac{1}{\mathcal{M}_t} - \frac{1}{\mathcal{M}} \right) + \text{E}_t \left[\frac{\text{Q}_{t+1}}{\text{Q}_t} \frac{\text{Y}_{t+1}}{\text{Y}_t} \psi\, \pi_{t+1} (1 + \pi_{t+1}) \right]$$

- Actual markup below desired ($\mathcal{M}_t < \mathcal{M}$): the first term is positive $\Rightarrow$ firms raise prices $\Rightarrow$ inflation.
- Actual markup above desired: firms cut prices $\Rightarrow$ deflation.
- The larger the gap, the larger the price change.

Interpreting the nonlinear Phillips curve

$$\psi\,\pi_t(1+\pi_t) = (1-\varepsilon) + \varepsilon\,\mathrm{mc}_t + \mathrm{E}_t\left[\frac{Q_{t+1}}{Q_t}\,\frac{Y_{t+1}}{Y_t}\,\psi\,\pi_{t+1}(1+\pi_{t+1})\right]$$

- Steady state with zero inflation ($\bar{\pi} = 0$): $0 = (1-\varepsilon) + \varepsilon\,\overline{\mathrm{mc}}$, so $\overline{\mathrm{mc}} = (\varepsilon-1)/\varepsilon = 1/\mathcal{M}$: the actual markup equals the desired one. (With positive trend inflation, $\bar{\pi} > 0$, the left side is positive, so $\overline{\mathrm{mc}} > 1/\mathcal{M}$ and $\bar{\mathcal{M}} < \mathcal{M}$.)
- Higher marginal cost $\Rightarrow$ higher inflation: firms' markups are squeezed below $\mathcal{M}$, and they raise prices to restore them.
- Higher expected inflation $\Rightarrow$ higher inflation today: if prices will rise in the future, firms start raising prices now, to spread the adjustment.

The role of ψ: how sticky prices are

- Divide by ψ: the effect of marginal cost on inflation is scaled by $1/\psi$.
- Higher $\psi \Rightarrow$ inflation reacts less to marginal cost. Changing prices is expensive, so firms adjust by less each period and spread the adjustment over time.
- $\psi = 0$: no cost of changing prices. Setting $\psi = 0$ in the equation gives $\mathrm{mc}_t = (\varepsilon-1)/\varepsilon$ in every period — the flexible-price markup rule.

Linearizing the Phillips curve (1/2): terms (1) and (2)

Around the zero-inflation steady state ($\bar{\pi} = 0$, $\overline{mc} = \frac{\varepsilon - 1}{\varepsilon}$, Y constant):

$$\underbrace{\psi\, \pi_t (1 + \pi_t)}_{(1)} = \underbrace{(1 - \varepsilon) + \varepsilon\, mc_t}_{(2)} + E_t \Big[\underbrace{\frac{Q_{t+1}}{Q_t} \frac{Y_{t+1}}{Y_t}\, \psi\, \pi_{t+1} (1 + \pi_{t+1})}_{(3)} \Big]$$

First-order Taylor expansion around the steady state $(\bar{x}, \bar{z})$

$$f(x) \approx f(\bar{x}) + f'(\bar{x})(x - \bar{x}) \qquad f(x, z) \approx f(\bar{x}, \bar{z}) + f_x(\bar{x}, \bar{z})(x - \bar{x}) + f_z(\bar{x}, \bar{z})(z - \bar{z})$$

1. $f(\pi_t) = \psi\, \pi_t (1 + \pi_t)$, with $f'(\pi_t) = \psi(1 + 2\pi_t)$. At $\bar{\pi} = 0$: $f(0) = 0$, $f'(0) = \psi$.
   $\Rightarrow$ $(1) \approx f(0) + f'(0)(\pi_t - 0) = \psi\, \pi_t$.
2. $g(mc_t) = (1 - \varepsilon) + \varepsilon\, mc_t$: already linear. In steady state, $g(\overline{mc}) = (1 - \varepsilon) + \varepsilon\, \overline{mc} = 0$. Use the percentage deviation $\widehat{mc}_t \equiv \frac{mc_t - \overline{mc}}{\overline{mc}}$ (to first order, the same as the log deviation used for the household).
   $\Rightarrow$ $(2) = \underbrace{(1 - \varepsilon) + \varepsilon\, mc_t}_{g(mc_t)} - \underbrace{\big[(1 - \varepsilon) + \varepsilon\, \overline{mc}\big]}_{=0} = \varepsilon\, (mc_t - \overline{mc}) = \varepsilon\, \overline{mc}\, \widehat{mc}_t = (\varepsilon - 1)\, \widehat{mc}_t$.

Result (after term (3), next slide)

$$\psi\, \pi_t = (\varepsilon - 1)\, \widehat{mc}_t + \beta \psi\, E_t \pi_{t+1}$$

Linearizing the Phillips curve (2/2): term (3)

$$\underbrace{\psi\,\pi_t(1+\pi_t)}_{(1)} = \underbrace{(1-\varepsilon)+\varepsilon\,mc_t}_{(2)} + E_t\Big[\underbrace{\tfrac{Q_{t+1}}{Q_t}\,\tfrac{Y_{t+1}}{Y_t}\,\psi\,\pi_{t+1}(1+\pi_{t+1})}_{(3)}\Big]$$

- The discount factor Q_{t+1}/Q_t. Firms belong to households, so they discount future profits as households do: $\frac{Q_{t+1}}{Q_t} = \beta\left(\frac{C_{t+1}}{C_t}\right)^{-\sigma}$. In steady state consumption is constant, so $\frac{Q_{t+1}}{Q_t} = \beta$.
- With Y also constant, the first two factors together, $z_t \equiv \frac{Q_{t+1}}{Q_t}\frac{Y_{t+1}}{Y_t}$, equal $\bar{z} = \beta$ in steady state.

Linearize $h(z_t, \pi_{t+1}) = z_t\,\psi\,\pi_{t+1}(1+\pi_{t+1})$ around $(\bar{z}, \bar{\pi}) = (\beta, 0)$:

- Derivatives: $h_z = \psi\,\pi_{t+1}(1+\pi_{t+1})$, $h_\pi = z_t\,\psi(1+2\pi_{t+1})$.
- At $(\beta, 0)$: $h = 0$, $h_z = 0$, $h_\pi = \beta\psi$.

$$(3) \approx h(\beta,0) + h_z(\beta,0)\,(z_t-\beta) + h_\pi(\beta,0)\,(\pi_{t+1}-0) = 0 + 0\cdot(z_t-\beta) + \beta\psi\,\pi_{t+1} = \beta\psi\,\pi_{t+1}$$

Deviations of z_t drop out: their coefficient is inflation, which is zero in steady state. So only β matters.

Result: take E_t and put the three terms together

$$\psi\,\pi_t = (\varepsilon-1)\,\widehat{mc}_t + \beta\psi\,E_t\pi_{t+1}$$

The linear Phillips curve in marginal cost

Collect the terms. Put (1), (2) and (3) back into the Phillips curve and take E_t of (3):

$$\psi\,\pi_t = (\varepsilon - 1)\,\widehat{\mathrm{mc}}_t + \beta\psi\,\mathrm{E}_t\pi_{t+1}$$

Divide by ψ:

Linear Phillips curve

$$\pi_t = \beta\,\mathrm{E}_t\pi_{t+1} + \frac{\varepsilon - 1}{\psi}\,\widehat{\mathrm{mc}}_t$$

- Inflation today depends on expected inflation tomorrow and on marginal cost today.
- Stickier prices (larger ψ): inflation reacts less to marginal cost.
- More competition (larger ε): inflation reacts more.
- β in front of expected inflation: future price changes are discounted.

Remaining problem

Marginal cost is not observed. Next: express it in terms of output.

Output gap

Marginal cost in terms of output: the ingredients

Marginal cost is not observed. We need four conditions, each in logs and in deviations from steady state (lowercase = log deviation, as before; the exception is mc_t, a level, whose log deviation is $\widehat{mc}_t$):

1. Marginal cost (firms): $mc_t = \frac{W_t/P_t}{A_t}$. Take logs: $\ln mc_t = \ln \frac{W_t}{P_t} - \ln A_t$. Subtract the steady state:

$$\underbrace{\ln mc_t - \ln \overline{mc}}_{\widehat{mc}_t} = \underbrace{\left(\ln \frac{W_t}{P_t} - \ln \frac{\overline{W}}{P}\right)}_{w_t} - \underbrace{\left(\ln A_t - \ln \bar{A}\right)}_{a_t} \quad \Longrightarrow \quad \widehat{mc}_t = w_t - a_t$$

2. Labor supply (households, derived earlier): $w_t = \sigma c_t + \varphi n_t$
3. Labor market: $Y_t = A_t N_t$. Take logs, $\ln Y_t = \ln A_t + \ln N_t$, and subtract the steady state:

$$y_t = a_t + n_t \quad \Longrightarrow \quad n_t = y_t - a_t$$

4. Goods market: adding the budget constraints (bonds in zero net supply, profits returned as D_t) gives $C_t = \left(1 - \frac{\psi}{2}\pi_t^2\right) Y_t$. Take logs and subtract the steady state ($\bar{\pi} = 0$, so $\ln \bar{C} = \ln \bar{Y}$):

$$c_t = y_t + \ln\left(1 - \frac{\psi}{2}\pi_t^2\right)$$

Linearize the last term around $\bar{\pi} = 0$: $f(\pi_t) = \ln\left(1 - \frac{\psi}{2}\pi_t^2\right)$, $f(0) = 0$, $f'(\pi_t) = \frac{-\psi\pi_t}{1 - \frac{\psi}{2}\pi_t^2}$, so $f'(0) = 0$. Hence $f(\pi_t) \approx 0$ and

$$c_t = y_t$$

The adjustment cost depends on π_t^2, so it does not survive a first-order approximation.

Marginal cost in terms of output: combining

Substitute step by step:

- Goods market ($c_t = y_t$) and labor market ($n_t = y_t - a_t$) into labor supply:

$$w_t = \sigma y_t + \varphi(y_t - a_t)$$

- This into marginal cost, $\widehat{mc}_t = w_t - a_t$:

$$\widehat{mc}_t = \sigma y_t + \varphi(y_t - a_t) - a_t = (\sigma + \varphi)\, y_t - (1 + \varphi)\, a_t$$

- Higher output $\Rightarrow$ more hours and more consumption, both requiring higher wages $\Rightarrow$ higher marginal cost.
- Higher productivity $\Rightarrow$ fewer hours per unit $\Rightarrow$ lower marginal cost.

This is where $\sigma + \varphi$ comes from: how fast marginal cost rises with output

- Larger φ (less elastic labor supply): more hours require a much higher wage.
- Larger σ (stronger wealth effect): richer households value leisure more and require a higher wage.
- Together with price stickiness, $\sigma + \varphi$ determines the slope of the Phillips curve.

Natural output and the output gap

Natural output y_t^n: output with flexible prices. Firms then charge the desired markup, so $\widehat{mc}_t = 0$:

$$0 = (\sigma + \varphi)\, y_t^n - (1 + \varphi)\, a_t \quad \Longrightarrow \quad y_t^n = \frac{1 + \varphi}{\sigma + \varphi}\, a_t$$

It moves with productivity, not with monetary policy. With $\sigma = 1$, $y_t^n = a_t$: hours do not respond, as the higher wage and the wealth effect cancel. With $\sigma > 1$ the wealth effect dominates and y_t^n rises by less than a_t.

Marginal cost is proportional to the output gap

$$\widehat{mc}_t = (\sigma + \varphi)\,(y_t - y_t^n) = (\sigma + \varphi)\, x_t, \qquad x_t \equiv y_t - y_t^n$$

- $x_t > 0$: overheating; costs above normal, firms want to raise prices. $x_t < 0$: slack.
- Not every fall in output is a policy problem: if productivity falls, y_t^n falls too.

In practice the output gap is estimated, e.g. by the US Congressional Budget Office (CBO) or the European Commission. These estimates are close to, but not the same as, the model's concept.

Equation 1: the New Keynesian Phillips curve

Substitute $\widehat{mc}_t = (\sigma + \varphi)x_t$ into the linear Phillips curve $\pi_t = \beta E_t \pi_{t+1} + \frac{\varepsilon - 1}{\psi} \widehat{mc}_t$:

NKPC

$$\pi_t = \beta \, E_t \pi_{t+1} + \kappa \, x_t + u_t, \qquad \kappa \equiv \frac{(\varepsilon - 1)(\sigma + \varphi)}{\psi}$$

- κ: the slope. Stickier prices (larger ψ) $\Rightarrow$ flatter; larger σ, φ or $\varepsilon \Rightarrow$ steeper.
- u_t: a cost-push shock. It is not in our derivation: with constant ε, $u_t = 0$. We add it because desired markups can vary: if ε_t varies, the desired markup deviates by μ_t, and the Phillips curve gains a term $\frac{\varepsilon - 1}{\psi} \mu_t \equiv u_t$. Wage markups and energy costs act similarly. It raises inflation even with a zero output gap.
- Iterating forward:

$$\pi_t = \kappa \sum_{k \geq 0} \beta^k E_t x_{t+k} + \sum_{k \geq 0} \beta^k E_t u_{t+k}$$

Inflation today is driven by expected future gaps and cost pressures.

Comparing with the traditional Phillips curve

Traditional (backward-looking)

$$\pi_t = \pi_{t-1} + \kappa x_t$$

- Inflation has inertia: it depends on past inflation.
- Disinflation requires a period of low output.

New Keynesian (forward-looking)

$$\pi_t = \beta E_t \pi_{t+1} + \kappa x_t$$

- Inflation depends on expected inflation.
- A credible central bank can lower inflation by changing expectations.

- The NK version is derived from firms' decisions; the traditional one was an empirical regularity.
- In the data, inflation shows more persistence than the pure NK version implies; empirical work often adds lagged inflation (a "hybrid" Phillips curve).

Galí and Gertler (1999), "Inflation Dynamics: A Structural Econometric Analysis," JME.

Equation 2: the dynamic IS curve (1/2)

Start from the linearized Euler equation and impose goods market clearing, $c_t = y_t$:

$$y_t = E_t y_{t+1} - \frac{1}{\sigma}(i_t - E_t \pi_{t+1} - \rho)$$

The same equation holds under flexible prices, with natural output y_t^n and the real rate that would prevail then, the natural rate r_t^n:

$$y_t^n = E_t y_{t+1}^n - \frac{1}{\sigma}(r_t^n - \rho)$$

Solve for r_t^n:

$$r_t^n = \rho + \sigma E_t \Delta y_{t+1}^n$$

- The natural rate is high when natural output is expected to grow: households expect to be richer and want to borrow, so a higher rate is needed to keep demand in line.

Equation 2: the dynamic IS curve (2/2)

Subtract the flexible-price equation from the actual one, using $x_t = y_t - y_t^n$:

Dynamic IS curve

$$x_t = E_t x_{t+1} - \frac{1}{\sigma}\left(i_t - E_t \pi_{t+1} - r_t^n\right)$$

- The output gap depends on the gap between the actual real rate and the natural rate.
- It is "IS" because it is the demand side: a higher real rate lowers spending, as in the IS–LM model.
- It is "dynamic" because expectations of the future gap enter.

Interpreting the IS curve and the natural rate

$$x_t = E_t x_{t+1} - \frac{1}{\sigma}\left(\underbrace{i_t - E_t \pi_{t+1}}_{\text{real rate}} - \underbrace{r_t^n}_{\text{natural rate}}\right)$$

- r_t^n: the natural rate of interest — the real rate that would keep output at its natural level.
- Policy is tight when the real rate is above r_t^n (the gap falls), and loose when it is below.
- With $y_t^n = \frac{1+\varphi}{\sigma+\varphi} a_t$:

  $$r_t^n = \rho + \sigma \frac{1+\varphi}{\sigma+\varphi} E_t \Delta a_{t+1}$$

  Expected productivity growth raises the natural rate.
- Demand shocks (e.g. to preferences or fiscal spending) also move r_t^n.
- Without shocks, $r_t^n = \rho$: the steady-state real rate.

Link to the Taylor rule

The intercept r^* in the Taylor rule is an estimate of the natural rate: the "neutral" real rate.

Equation 3: monetary policy

- Two equations, three unknowns: x_t, π_t, i_t.
- The model is closed by a description of monetary policy, e.g. a Taylor-type rule:

$$i_t = \rho + \phi_\pi \pi_t + \phi_x x_t$$

- Or by an optimal policy: the central bank chooses i_t to maximize welfare (next lecture).

Money in the background

The central bank's balance sheet adjusts so that i_t is achieved. In the long run the quantity-theory link holds: money growth equals inflation.

The three-equation New Keynesian model

Demand: dynamic IS curve

$$x_t = E_t x_{t+1} - \frac{1}{\sigma}\left(i_t - E_t \pi_{t+1} - r_t^n\right)$$

Supply: New Keynesian Phillips curve

$$\pi_t = \beta E_t \pi_{t+1} + \kappa x_t + u_t$$

Policy

$$i_t = \text{rule, or optimal choice}$$

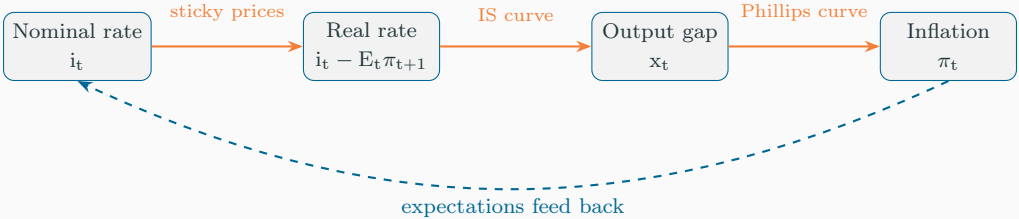
The parameters and where they appear

Parameter	Meaning	Typical value	Appears in
β	discount factor (patience)	0.99 (quarterly)	NKPC; $\rho = -\ln \beta$
σ	inverse intertemporal elasticity	1–2	IS slope $1/\sigma$; κ
φ	inverse Frisch elasticity	1–2	κ; natural output
ε	elasticity of demand (competition)	6 (20% markup)	κ
ψ	cost of changing prices	large	κ
κ	slope of the Phillips curve	0.01–0.1	NKPC

- $\kappa = (\varepsilon - 1)(\sigma + \varphi)/\psi$ combines preferences, competition and price stickiness.
- Estimates of κ are small: a 1% output gap moves quarterly inflation by a small amount. A flat Phillips curve is a recurrent finding, especially before 2021.
- Every parameter has an economic interpretation: this is the advantage of deriving the model rather than assuming it.

Mechanism

The transmission mechanism



- Because prices are sticky, a change in i_t changes the real rate.
- The real rate moves demand (IS); demand moves marginal cost and inflation (NKPC).
- Because both equations are forward-looking, expected future policy matters as much as today's rate.
- With $\psi = 0$: $x_t = 0$ always, and monetary policy only affects inflation.

Example: a monetary tightening, step by step

The central bank raises i_t by 1 percentage point, unexpectedly.

1. Real rate: prices are sticky, so expected inflation does not rise one-for-one. The real rate $i_t - E_t\pi_{t+1}$ rises.
2. Households (IS curve): saving becomes more attractive; consumption falls. Output falls below natural: $x_t < 0$.
3. Labor market: firms need fewer hours; the real wage falls, and so does marginal cost.
4. Firms (Phillips curve): with lower costs, firms raise prices by less. Inflation falls.
5. Expectations: if the tightening is expected to last, expected future gaps are negative too, and inflation falls further today.

The cost of disinflation

Inflation falls only because output falls below its natural level. How much output must be given up depends on κ: the flatter the Phillips curve, the higher the cost.

Two kinds of shocks

Shocks to r_t^n (demand, productivity)

- Move the natural rate.
- If the central bank sets the real rate equal to r_t^n, then $x_t = 0$ and $\pi_t = 0$.
- No trade-off.

Cost-push shocks u_t

- E.g. markup or energy cost shocks.
- Stabilizing inflation requires a negative output gap.
- A genuine trade-off.

The "divine coincidence"

Without cost-push shocks, stabilizing inflation also closes the output gap. Cost-push shocks break it — and that is where optimal policy, and the mandate, start to matter.

Blanchard and Galí (2007), JMCB.

Summary and next lecture

- Households give the Euler equation ⇒ the IS curve: demand depends on the expected path of real rates.
- Monopolistic competition gives price-setting firms with a markup; costly price adjustment makes pricing forward-looking ⇒ the Phillips curve.
- Marginal cost is proportional to the output gap; the natural rate r_t^n is the real rate that keeps output at its natural level.
- Sticky prices give the central bank control of the real rate.

Next Lecture

What should the central bank do? Long-run optimal inflation, short-run stabilization, and why the answer looks like a Taylor rule.

Readings: Galí (2015), Monetary Policy, Inflation, and the Business Cycle, ch. 3; Clarida, Galí and Gertler (1999), JEL, section 2; Rotemberg (1982), JPE.