Monetary Theory and Policy

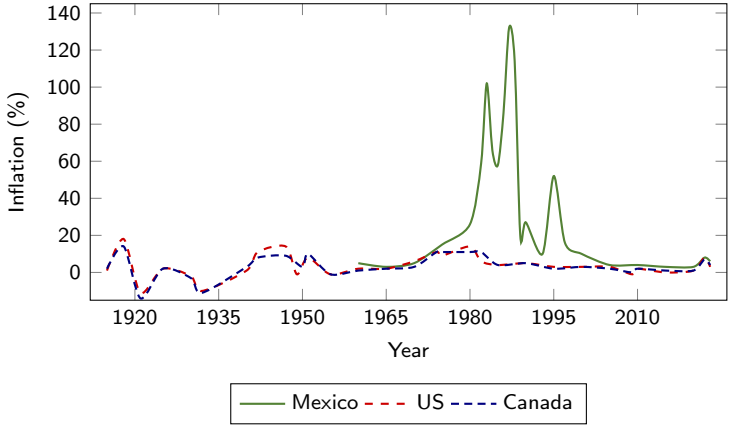
Alberto Ramírez de Aguilar

Universidad Anáhuac

Fall 2026

Inflation in North America: A Long View

- The figure below shows annual CPI inflation for the United States, Canada, and Mexico from 1915 to the present. The contrast is striking: while the US and Canada have maintained relatively moderate inflation for most of the period, Mexico experienced a severe **inflation crisis** in the 1980s, with annual rates exceeding **130%** at the peak in 1987.



Why Does Low and Stable Inflation Matter?

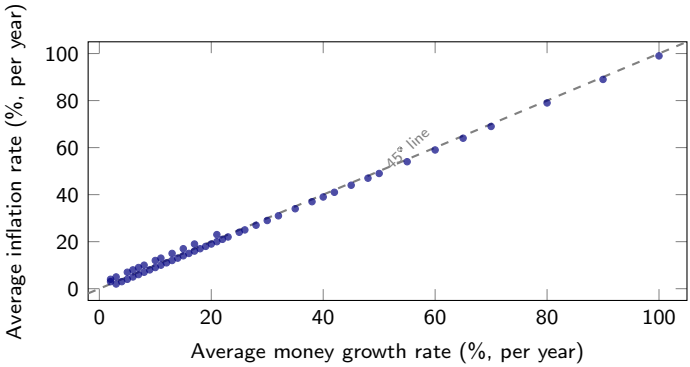
- **Low and stable inflation** is the primary mandate of most central banks. But why? What is actually harmful about inflation?

- **High inflation erodes the value of nominally denominated assets**:
  - Bank deposits and cash holdings lose purchasing power. A household holding $10,000 in a savings account during Mexico's 1987 inflation saw its real value fall by more than half in one year.
  - Real wages fall if nominal wages are not perfectly indexed. Workers whose contracts were set at the beginning of the year suffer a real pay cut by year's end.

- **Unpredictable inflation arbitrarily redistributes wealth and creates uncertainty**:
  - Firms and households cannot form reliable expectations about future prices - how much wage increase should workers bargain for? What is the real value of a long-term investment? Economic decisions become distorted.

Money and Prices: The Long-Run Evidence

- The most robust empirical fact in monetary economics: **in the long run, money growth and inflation move together, nearly one-for-one.** McCandless and Weber (1995) studied 110 countries over 30 years and found a correlation between money growth and inflation close to 1.



- **In the long run, there is no trade-off** - faster money creation does not buy permanently higher output. It buys only higher prices.

Money and Real Variables: The Short Run vs. the Long Run

- The long-run picture is clear: money growth drives inflation, not output. But **the short run is different**. A large body of evidence - from Friedman and Schwartz (1963) to Romer and Romer (2004) and Christiano, Eichenbaum, and Evans (1999) - shows that monetary shocks have significant and persistent effects on real output over horizons of one to three years.

- **The canonical evidence**: following a contractionary monetary policy shock, real GDP falls, employment drops, and investment declines. Prices respond more slowly - the "price puzzle" - and return to a lower level eventually. In the long run, output returns to trend: **money is neutral in the long run but not in the short run**.

- **Why does this short-run non-neutrality arise?** Three leading explanations, each with a formal model in this course:
  - **Imperfect information** (Lucas 1972).
  - **Sticky prices** (New Keynesian models).
  - **Financial frictions**.

What Does a Central Bank Do?

- Every modern economy has a central bank charged with managing the supply of money and credit. Their **legal mandates** vary somewhat across countries, but converge on the same core objectives:
  - *"... to maintain a low and stable inflation."* **Banco de México**
  - *"... to keep inflation low, stable and predictable."* **Bank of Canada**
  - *"... to promote maximum employment and stable prices."* **Federal Reserve**
- Price stability and maximum employment are not chosen arbitrarily. They reflect the two most costly macroeconomic failures that monetary policy can directly address. This course is about understanding *why* these are the right goals, *how* central banks can achieve them, and *what* can go wrong when they fail.
- But before building models, we should look at the data. How well have central banks actually performed? What does inflation look like across time and countries? The next slides show the evidence that motivates everything in this course.

Central Banking: Science, Art, and the Role of Models

- Some describe central banking as *"a complex blend of art and science - in effect, it is an exercise in risk management."* We will focus on the **science**: the formal models and empirical tools that discipline our thinking.

- **Theoretical models** serve two distinct purposes:
  - **Explain a mechanism**: why do banks exist? Diamond (1984): banks as delegated monitors. What makes an object suitable as a medium of exchange? Kiyotaki and Wright (1989): model of endogenous money.
  - **Make a counterintuitive point**: Lucas (1972) shows that a short-run trade-off between unemployment and inflation can exist in the data and yet not be exploitable by policy. Models prevent us from reasoning by analogy from misleading empirical patterns.

- **Our approach in this course**: we build formal models with explicit microfoundations, derive their quantitative predictions, and confront them with the historical and cross-country evidence. Theory and evidence in constant dialogue.

The Challenge: Why Is Money Hard to Model?

- Despite its central importance in economic life, **money is surprisingly difficult to model**. The fundamental challenge is that money is a **dominated asset**: it pays zero nominal interest, while government bonds pay a positive return. Yet people voluntarily hold money. Why?

- This is known as the **rate-of-return dominance puzzle**. In a frictionless Arrow-Debreu general equilibrium model, no one would voluntarily hold a dominated asset - money would have no role. This is essentially what the Classical Growth Model shows: a complete, frictionless economy has no need for money.

- To give money a role, we need to introduce a friction. The literature has proposed several:
  - **Put money directly in utility** (MIU): assume agents value real balances. Tractable, but ad hoc.
  - **Cash-in-advance constraint**: require cash for transactions. Microfounded liquidity need, but mechanical.
  - **Search and matching** (Kiyotaki-Wright): bilateral trade with anonymity makes money essential. Rich microfoundations, more complex.

Outline

1. A Primer on Monetary Economics
2. Foundations of Money
3. The Welfare Cost of Inflation
4. Deeper Foundations of Money
5. Quantity Theory, Expectations, and Policy
6. Time Inconsistency and Central Bank Design
7. Monetary-Fiscal Interactions
8. The Role of Private Banks

The Classical Growth Model

- This is a course on **monetary theory and policy** - so, naturally, we begin with a model that has no money in it.

- The **Classical (Neoclassical) Growth Model** - built on the work of Ramsey (1928), Cass (1965), and Koopmans (1965) - is the workhorse of modern macroeconomics. In this model, households optimize over an infinite horizon, firms rent capital and labor in competitive markets, and the economy converges to a well-defined steady state.

- We study it for three reasons:
  - Provides the **benchmark**: a fully specified general equilibrium economy against which we will compare every monetary model in this course.
  - Matches the **Kaldor facts**: a set of robust empirical regularities about long-run growth - which gives us confidence that the model captures something real before we start modifying it.
  - **Neutrality result**: money, if introduced, would have no real effects. Prices adjust, output does not. This is the puzzle that motivates everything that follows.

Kaldor's Facts (1957)

- Nicholas Kaldor pointed out in a now seminal article, *A Model of Economic Growth*, six facts about economic growth. Since then, these have become known in the literature as **stylized facts**:

  1. Output per worker has grown at a roughly constant rate.
  2. Capital per worker grows over time at basically the same rate.
  3. The Capital/Output ratio has been constant over time.
  4. The return of Capital has been constant.
  5. The share of Capital and Labor expenditures used in production has been constant.
  6. Real wages have increased through time.

- The Classical Growth Model, which we now study, has become fairly popular since it accounts for all these facts.

GDP Evolution: United States and Mexico

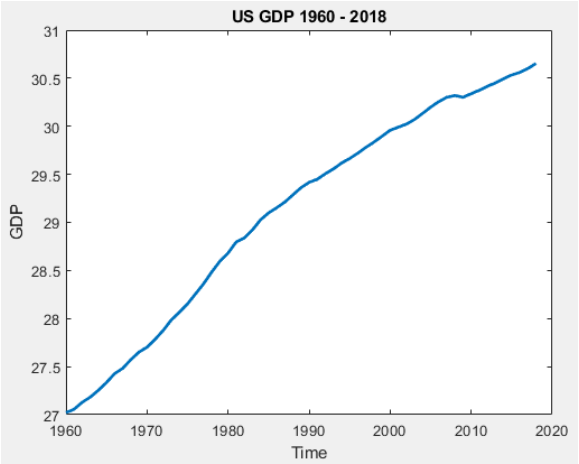


Figure: US GDP Since 1960.

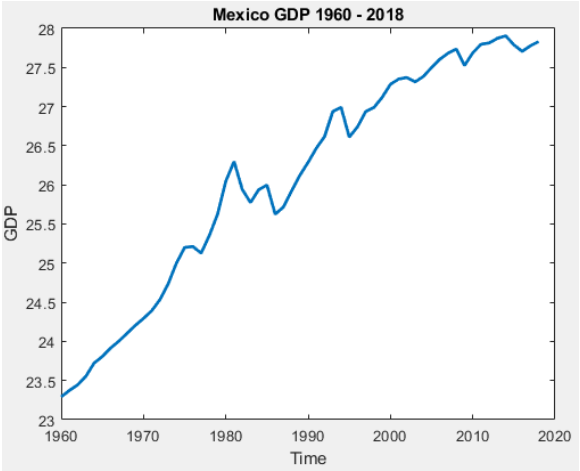


Figure: Mexico's GDP Since 1960.

The Classical Growth Model

- The basic ingredients of the Classical Growth Model are the following:
  - Households: They seek a consumption allocation c to maximize their utility. They provide labor n_t and are owners of capital k_t, receiving from both an income each period. Finally, households do investment decisions on future capital.
  - Firms: Profit maximizers who rent both labor and capital to households. They have a constant returns to scale technology.
  - Information: There is no risk in this economy, all sequences are deterministic.
  - Competitive Equilibrium Concept: Both household and firms **cannot** influence prices, they take them as given.

The Classical Growth Model: Firms

- This model is populated with a mass of identical firms. Each one faces the following production function:

$$Y_t = F(K_t, N_t) = AK_t^{\alpha} N_t^{1-\alpha},$$

  which is a Cobb-Douglas production technology.

- Firms must pay a wage w_t for each unit of labor they hire and a rent R_t for each unit of capital utilized. We normalize the price of the consumption good to one.

- Hence, firms solve the following problem:

$$\max_{\{N_t, K_t\}} AK_t^{\alpha} N_t^{1-\alpha} - w_t N_t - R_t K_t \quad \text{subject to}$$

$$N_t \geq 0,$$

$$K_t \geq 0.$$

The Classical Growth Model: Firms

- The FOCs of this problem imply:

$$w_t = F_N(K_t, N_t) = (1-\alpha)A\left[\frac{K_t}{N_t}\right]^{\alpha},$$

$$R_t = F_K(K_t, N_t) = \alpha A\left[\frac{N_t}{K_t}\right]^{1-\alpha}.$$

- Notice that the share of a firm's expenses in labor is given then by:

$$w_t N_t = (1-\alpha)A\left[\frac{K_t}{N_t}\right]^{\alpha} N_t = (1-\alpha)AK_t^{\alpha}N_t^{1-\alpha} = (1-\alpha)Y_t.$$

- Hence, this implies that the firm's expenses in labor relative to output is constant:

$$\frac{w_t N_t}{Y_t} = 1-\alpha,$$

which is a Kaldor fact! Similarly for capital.

The Classical Growth Model: Households

- Households seek to maximize their utility, which depends on consumption and labor. They own the capital stock of the economy and decide how much to invest considering depreciation.
- Labor is perfectly movable within a period. However, capital is not. Households then need to decide the stock of capital they wish to have in $t+1$ during period t. They take as given the initial capital stock k_0.
- They solve the following maximization problem:

$$\max_{\{c_t, n_t, k_{t+1}, i_t\}} \sum_{t=0}^{\infty} \beta^t \left[log(c_t) - \eta \frac{n_t^{1+\nu}}{1+\nu} \right] \quad \text{subject to}$$

$$c_t + i_t = w_t n_t + R_t k_t,$$

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

$$c_t, \quad i_t, \quad n_t, \quad k_{t+1} \geq 0,$$

$$k_0 \quad \text{given.}$$

The Classical Growth Model: Households

- This problem can be rewritten as follows:

$$\max_{\{c_t, n_t, k_{t+1}\}} \sum_{t=0}^{\infty} \beta^t \left[log(c_t) - \eta \frac{n_t^{1+\nu}}{1+\nu} \right] \quad \text{subject to}$$

$$c_t + k_{t+1} = w_t n_t + R_t k_t + (1-\delta) k_t$$

$$c_t, \quad k_{t+1} \geq 0,$$

$$k_0 \quad \text{given},$$

for which the Lagrangian is given by:

$$\mathcal{L} = \sum_{t=0}^{\infty} \beta^t \left[log(c_t) - \eta \frac{n_t^{1+\nu}}{1+\nu} \right] + \sum_{t=0}^{\infty} \lambda_t [w_t n_t + R_t k_t + (1-\delta) k_t - c_t - k_{t+1}]$$

$$= ... + \lambda_t [w_t n_t + R_t k_t + (1-\delta) k_t - c_t - k_{t+1}]$$

$$+ \lambda_{t+1} [w_{t+1} n_{t+1} + R_{t+1} k_{t+1} + (1-\delta) k_{t+1} - c_{t+1} - k_{t+2}] + ...$$

The Classical Growth Model: Households

- The FOCs of the household are:

$$\frac{\partial \mathcal{L}}{\partial c_t} = 0 \Rightarrow \frac{\beta^t}{c_t} = \lambda_t,$$
$$\frac{\partial \mathcal{L}}{\partial n_t} = 0 \Rightarrow \eta \beta^t n_t^{\nu} = w_t \lambda_t,$$
$$\frac{\partial \mathcal{L}}{\partial k_{t+1}} = 0 \Rightarrow \lambda_t = [R_{t+1} + 1 - \delta]\lambda_{t+1}.$$

- These FOCs imply both the intra-temporal equation:

$$\eta c_t n_t^{\nu} = w_t,$$

and the Euler equation (inter-temporal equation):

$$\frac{c_{t+1}}{\beta c_t} = R_{t+1} + 1 - \delta.$$

The Classical Growth Model: Competitive Equilibrium

- In order to close up the model, we need to specify a notion of equilibrium.
- A **competitive equilibrium** are prices $\{\hat{w}_t, \hat{R}_t\}_{t=0}^{\infty}$, an allocation for the household $\{\hat{c}_t, \hat{n}_t, \hat{k}_{t+1}\}_{t=0}^{\infty}$, and a demand vector for firms $\{\hat{Y}_t, \hat{N}_t, \hat{K}_{t+1}\}_{t=0}^{\infty}$ such that:
  1. Given prices $\{\hat{w}_t, \hat{R}_t\}_{t=0}^{\infty}$, the household's allocation $\{\hat{c}_t, \hat{n}_t, \hat{k}_{t+1}\}_{t=0}^{\infty}$ maximizes their utility subject to their period by period budget constraint.
  2. Given prices $\{\hat{w}_t, \hat{R}_t\}_{t=0}^{\infty}$, the firm's allocation $\{\hat{Y}_t, \hat{N}_t, \hat{K}_{t+1}\}_{t=0}^{\infty}$ maximizes its profits.
  3. Markets clear every period:

$$\hat{c}_t + \hat{k}_{t+1} = \hat{Y}_t + (1-\delta)\hat{k}_t,$$
$$\hat{n}_t = \hat{N}_t,$$
$$\hat{k}_{t+1} = \hat{K}_{t+1}.$$

Solving for an Equilibrium

- To attain the equilibrium, we need to use all the FOCs we have computed before and combine them with market clearing conditions.

- Notice we have (basically) three unknowns every period (c_t, n_t, k_{t+1}), so we need to attain three equations to solve for these variables. These equations are given by (make sure you understand where they come from):

$$\eta c_t n_t^{\nu} = (1-\alpha) A \left[\frac{k_t}{n_t}\right]^{\alpha},$$

$$\frac{c_{t+1}}{\beta c_t} = \alpha A \left[\frac{n_{t+1}}{k_{t+1}}\right]^{1-\alpha} + 1 - \delta,$$

$$c_t + k_{t+1} = A k_t^{\alpha} n_t^{1-\alpha} + (1-\delta) k_t.$$

- Problem: These equations are non-linear and we cannot get a closed form solution (but we can use a software like MATLAB to compute the solution).

Economic Intuition: Reading the Equilibrium

- What are these conditions really telling us?
- **Intratemporal:** substitution between consumption and leisure = real wage = marginal product of labor. More capital per worker $\Rightarrow$ higher wage.
- **Intertemporal:** the Euler equation - households give up consumption today only if compensated by the net return to capital. Diminishing returns mean this compensation falls as k_t rises.
- **Kaldor's Facts:** Cobb-Douglas $\Rightarrow$ constant factor shares $\alpha, 1-\alpha$ (*stable income shares*); the Euler equation pins a constant $r^\star$ (*stable return to capital*).
- **Bottom line:** wages rise as capital deepens, while diminishing returns pull r down - together, these two forces are what will discipline how the economy evolves over time.

Steady State of the Economy

- The next best thing we can do is to solve for the economy's **Steady State**.
- An economy is at its steady state whenever **all** variables become constant, i.e.:

$$c_t = c^\star, \quad n_t = n^\star, \quad k_{t+1} = k^\star.$$

- In a steady state, the three equations described above simplify and become:

$$\eta c n^{\nu+\alpha} = (1-\alpha) A k^\alpha,$$

$$\frac{1}{\beta} = \alpha A \left[\frac{n}{k}\right]^{1-\alpha} + 1 - \delta,$$

$$c = A k^\alpha n^{1-\alpha} - \delta k.$$

Steady State of the Economy

- We can solve analytically for the steady state. Actually we only need to find the value of steady state capital (since we can then use the other equations to find c, n).
- If we define $\tilde{A} = \alpha A[(1-\alpha)A]^{\frac{1}{\nu+\alpha}}$ and $\gamma = \frac{\alpha}{\nu+\alpha}$, it can be shown (please verify it) that:

$$k^{\star} = \left[\frac{\beta\alpha\tilde{A}}{1+\beta(1-\delta)}\right]^{\frac{1}{(1-\gamma)(1-\alpha)}}.$$

- What happens to the steady state capital whenever parameters change?

Converging to the Steady State

- If we are in a situation were $k_0 \neq k^\star$, it can be shown that over the course of time the economy will converge to its steady state.

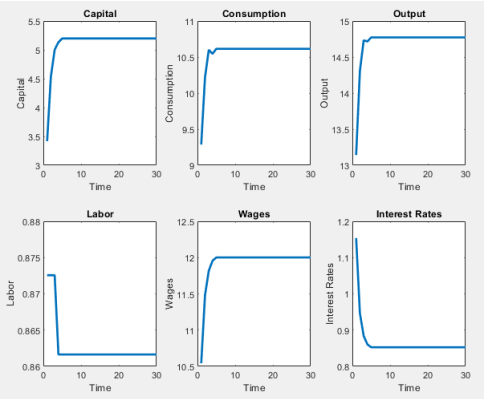


Figure: Convergence to the Steady State.

Economic Intuition: Steady State & Transition

- A steady state is not an economy that "stopped" - saving exactly offsets depreciation, so net investment is zero.
- **Why it converges:** diminishing returns drive the transition. If $k_0 < k^\star$, capital is scarce, its return is high, and it accumulates fast - then *slows* as $k_t \to k^\star$.
- Mirror image if $k_0 > k^\star$: low returns $\Rightarrow$ decumulation back down to $k^\star$. Either way, the economy is pulled to the same resting point.
- Same logic as **conditional convergence**: economies below their own steady state grow faster, simply because diminishing returns raise their investment incentive.
- **Kaldor's Facts:** at $k^\star$, the capital-output ratio $k^\star/y^\star$ is constant - Kaldor's *stable capital-to-output ratio*.
- **Takeaway:** diminishing returns to capital is the one mechanism doing all the work - it makes the steady state unique, stable, and self-correcting (the same force behind convergence in the Solow model).

Social Planner's Problem

- We now seek to answer the following questions: is the competitive equilibrium outcome efficient? Given its technology constraints, can the economy do better?

- In order to answer both questions, we now turn to analyze what is known in the literature as the **Social Planner's Problem** (SPP).
  - The SPP captures the idea of a **centralized** economy, where a benevolent outside agent (the Social Planner) chooses/dictates the allocations that both firms and households must consume.
  - The social planner seeks to maximize welfare in the economy subject to the technology constraints that are in the economy.
  - In the context of the Classical Model, the SPP becomes a maximization problem seeking to give the household the greatest utility, subject to the feasibility constraint of the economy (that is consumption plus investment cannot exceed output).
  - **The social planner does not consider prices when optimizing, she ONLY considers feasibility**.

Social Planner's Problem

- The SPP is then:

$$\max_{\{c_t, n_t, k_{t+1}\}} \sum_{t=0}^{\infty} \beta^t \left[log(c_t) - \eta \frac{n_t^{1+\nu}}{1+\nu} \right] \quad \text{subject to}$$

$$c_t + k_{t+1} = A k_t^{\alpha} n_t^{1-\alpha} + (1-\delta) k_t,$$

$$c_t, \quad n_t, \quad k_{t+1} \geq 0$$

$$k_0 \quad \text{given.}$$

Social Planner's Problem and the First Welfare Theorem

- The FOCs of the SPP are given by:

$$\frac{\beta^t}{c_t} = \lambda_t,$$

$$\eta \beta^t n_t^{\nu} = (1-\alpha) A \left[\frac{k_t}{n_t}\right]^{\alpha} \lambda_t,$$

$$\lambda_t = \left[\alpha A \left[\frac{n_{t+1}}{k_{t+1}}\right]^{1-\alpha} + 1 - \delta\right] \lambda_{t+1}.$$

- Notice that these equations yield exactly the same three equations (once we get rid of the λ's) that must hold in a competitive equilibrium!

- Hence, the **First Welfare Theorem** holds in the Classical Growth Model: any allocation that results from a competitive equilibrium is efficient.

Why Money Does Not Appear in the Classical Model

- The Classical Growth Model is a **frictionless** economy: goods, labor, and capital are traded costlessly and instantaneously. There is no need for a medium of exchange.

- As a consequence, money has no role in the model. Any household that chose to hold money rather than capital would earn a lower return - money is **dominated** as a store of value. In equilibrium, no one would willingly hold it.

- To give money a meaningful role, we must add a reason to hold it. Three main approaches appear in the literature:
  - **Money in the Utility Function (MIU)**: money yields utility directly, capturing in reduced form the transaction services it provides.
  - **Cash-in-advance (CIA)**: households must hold sufficient money balances before they can purchase consumption goods.
  - **Search and matching**: money arises endogenously as a solution to the double coincidence of wants problem in decentralized trade.

- We begin with the MIU approach, due to Sidrauski (1967), as it is tractable and delivers sharp, interpretable results.

Money in the Utility Function: Overview

- The key idea of the MIU approach is to place real money balances $m_t \equiv M_t/P_t$ directly into the household's period utility function, where M_t is the nominal money stock and P_t is the aggregate price level.

- **Justification**: money reduces the time and effort needed to complete transactions. Rather than modeling these frictions explicitly, we assume they generate utility directly. This is a reduced-form shortcut.

- **Criticism**: the MIU approach does not explain *why* money is useful - it simply assumes it is. We will later study models that derive the role of money from first principles.

- The MIU model modifies the Classical Growth Model in exactly two ways:
  1. Real money balances m_t enter the household's utility function.
  2. A government exists that controls the nominal money supply and transfers seigniorage revenue lump-sum to households.

MIU: The Government

- The government commits to a constant nominal money growth rate $\mu \geq -1$:

$$M_{t+1} = (1 + \mu) M_t.$$

- All newly issued money is transferred to households as a lump-sum nominal transfer $T_t = M_{t+1} - M_t$. The real value of this transfer is:

$$\tau_t \equiv \frac{T_t}{P_t} = \frac{M_{t+1} - M_t}{P_t} = m_{t+1}(1 + \pi_{t+1}) - m_t,$$

  where $\pi_{t+1} \equiv (P_{t+1} - P_t)/P_t$ is the net inflation rate between t and $t + 1$.

- In equilibrium, the price level adjusts so that the money market clears: $M_t = P_t m_t$ for all t.

- In a **steady state**, real balances are constant, so $(1 + \mu)P_t = P_{t+1}$, which implies:

$$\pi^\star = \mu.$$

  The steady-state inflation rate equals the money growth rate.

MIU: Households

- The stock of real balances $m_t \equiv M_t/P_t$ is chosen at the end of period $t-1$ and provides liquidity services during period t.

- The household's problem is:

$$\max_{\{c_t,\, n_t,\, k_{t+1},\, m_{t+1}\}} \sum_{t=0}^{\infty} \beta^t \left[\log(c_t) + \chi \log(m_t) - \eta \frac{n_t^{1+\nu}}{1+\nu}\right] \quad \text{subject to}$$

$$c_t + k_{t+1} + m_{t+1}(1+\pi_{t+1}) = w_t n_t + R_t k_t + (1-\delta)k_t + m_t + \tau_t,$$

$$c_t, \quad k_{t+1}, \quad m_{t+1} \geq 0, \qquad k_0, \quad m_0 \text{ given.}$$

- Note: the term $m_{t+1}(1+\pi_{t+1})$ is the period-t resource cost of carrying m_{t+1} real balances into period $t+1$. In nominal terms, acquiring M_{t+1} dollars at price P_t per unit of consumption costs $M_{t+1}/P_t = m_{t+1}(1+\pi_{t+1})$ units of consumption today.

MIU: Households

- The Lagrangian of the household's problem is:

$$\mathcal{L} = \sum_{t=0}^{\infty} \beta^t \left[\log(c_t) + \chi \log(m_t) - \eta \frac{n_t^{1+\nu}}{1+\nu} \right] +$$

$$\sum_{t=0}^{\infty} \lambda_t \left[w_t n_t + R_t k_t + (1-\delta) k_t + m_t + \tau_t - c_t - k_{t+1} - m_{t+1}(1+\pi_{t+1}) \right]$$

- The variable m_{t+1} appears as a cost in the period-t constraint and as a benefit in both period-$t+1$ utility and the period-$t+1$ constraint. Collecting those terms:

$$\begin{aligned} = \ & \ldots + \lambda_t [\ldots - m_{t+1}(1+\pi_{t+1})] \\ & + \beta^{t+1} \left[\log(c_{t+1}) + \chi \log(m_{t+1}) - \eta \frac{n_{t+1}^{1+\nu}}{1+\nu} \right] \\ & + \lambda_{t+1} [w_{t+1} n_{t+1} + R_{t+1} k_{t+1} + (1-\delta) k_{t+1} + m_{t+1} + \tau_{t+1} - c_{t+1} - k_{t+2} - m_{t+2}(1+\pi_{t+2})] + \end{aligned}$$

MIU: Households

- The FOCs of the household are:

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial c_t} = 0 &\Rightarrow \frac{\beta^t}{c_t} = \lambda_t, \\
\frac{\partial \mathcal{L}}{\partial n_t} = 0 &\Rightarrow \eta \beta^t n_t^{\nu} = w_t \lambda_t, \\
\frac{\partial \mathcal{L}}{\partial k_{t+1}} = 0 &\Rightarrow \lambda_t = \left[R_{t+1} + 1 - \delta\right] \lambda_{t+1}, \\
\frac{\partial \mathcal{L}}{\partial m_{t+1}} = 0 &\Rightarrow \lambda_t (1 + \pi_{t+1}) = \frac{\chi \beta^{t+1}}{m_{t+1}} + \lambda_{t+1}.
\end{aligned}$$

- The first three FOCs are **identical** to those of the Classical Growth Model. The fourth is new: it equates the opportunity cost of holding money (left-hand side) to its marginal utility plus its continuation value next period (right-hand side).

MIU: Households

- These FOCs imply the same intra-temporal and Euler equations as the Classical Model:

$$\eta c_t n_t^{\nu} = w_t, \qquad \frac{c_{t+1}}{\beta c_t} = R_{t+1} + 1 - \delta.$$

- The new FOC for m_{t+1} delivers a **money demand** condition. Substituting $\lambda_t = \beta^t / c_t$ and $\lambda_{t+1} = \beta^{t+1}/c_{t+1}$ into it and dividing by β^{t+1}:

$$\frac{1 + \pi_{t+1}}{\beta c_t} = \frac{\chi}{m_{t+1}} + \frac{1}{c_{t+1}}.$$

Using the Euler equation to substitute $1/(\beta c_t) = (R_{t+1} + 1 - \delta)/c_{t+1}$:

$$\frac{(1 + \pi_{t+1})(R_{t+1} + 1 - \delta) - 1}{c_{t+1}} = \frac{\chi}{m_{t+1}}.$$

Defining the **nominal interest rate** via the Fisher equation, $(1 + i_{t+1}) \equiv (R_{t+1} + 1 - \delta)(1 + \pi_{t+1})$, this simplifies to:

$$\frac{\chi \, c_{t+1}}{m_{t+1}} = i_{t+1}.$$

MIU: Competitive Equilibrium

- A **competitive equilibrium** in the MIU model consists of prices $\{\hat{w}_t, \hat{R}_t, \hat{P}_t\}_{t=0}^{\infty}$, a household allocation $\{\hat{c}_t, \hat{n}_t, \hat{k}_{t+1}, \hat{m}_{t+1}\}_{t=0}^{\infty}$, and firm demands $\{\hat{Y}_t, \hat{N}_t, \hat{K}_t\}_{t=0}^{\infty}$ such that:
  1. Given prices, the household allocation maximizes utility subject to the budget constraint.
  2. Given prices, the firm allocation maximizes profits.
  3. Goods, labor, capital, and money markets clear every period:

$$\hat{c}_t + \hat{k}_{t+1} + \hat{m}_{t+1}(1 + \hat{\pi}_{t+1}) = \hat{Y}_t + (1 - \delta)\hat{k}_t + \hat{m}_t + \tau_t,$$

$$\hat{n}_t = \hat{N}_t, \qquad \hat{k}_t = \hat{K}_t, \qquad M_t = \hat{P}_t \hat{m}_t.$$

- Substituting the government's budget $\tau_t = \hat{m}_{t+1}(1 + \hat{\pi}_{t+1}) - \hat{m}_t$ into the goods market clearing condition, the **money terms cancel**:

$$\hat{c}_t + \hat{k}_{t+1} = \hat{Y}_t + (1 - \delta)\hat{k}_t.$$

The consolidated resource constraint is **identical** to the Classical Growth Model.

Solving for an Equilibrium

- We now have four unknowns per period $(c_t, n_t, k_{t+1}, m_{t+1})$, requiring four equations. The first three - governing the real allocation - are **identical to those of the Classical Growth Model** (make sure you understand why):

$$\eta c_t n_t^{\nu} = (1-\alpha) A \left[\frac{k_t}{n_t}\right]^{\alpha},$$

$$\frac{c_{t+1}}{\beta c_t} = \alpha A \left[\frac{n_{t+1}}{k_{t+1}}\right]^{1-\alpha} + 1 - \delta,$$

$$c_t + k_{t+1} = A k_t^{\alpha} n_t^{1-\alpha} + (1-\delta) k_t.$$

- A fourth equation pins down real money balances, given the money growth rate μ set by the government:

$$\frac{\chi\, c_{t+1}}{m_{t+1}} = i_{t+1}, \qquad (1+i_{t+1}) = \left[\alpha A \left[\frac{n_{t+1}}{k_{t+1}}\right]^{1-\alpha} + 1 - \delta\right](1+\mu).$$

- The real block (c_t, n_t, k_{t+1}) can be solved entirely independently of money, and m_{t+1} is then determined residually.

Steady State of the MIU Economy

- In a steady state, all variables are constant: $c_t = c^\star$, $n_t = n^\star$, $k_t = k^\star$, $m_t = m^\star$, $\pi_t = \pi^\star = \mu$.
- The three equations for the real block are **identical** to those of the Classical Growth Model. Hence, the steady-state values of consumption, labor, and capital $(c^\star, n^\star, k^\star)$ are **the same as before** - they do not depend on μ.
- The nominal interest rate in steady state is determined by the Fisher equation:

$$1 + i^\star = \left(\frac{1}{\beta}\right)(1 + \mu),$$

since $R^\star + 1 - \delta = 1/\beta$ from the Euler equation.

- Steady-state real money balances are then pinned down by the money demand condition:

$$m^\star = \frac{\chi\, c^\star}{i^\star} = \frac{\chi\, c^\star \beta}{1 + \mu - \beta}.$$

Higher money growth μ raises the nominal interest rate, increasing the opportunity cost of holding money and **reducing** $m^\star$.

Neutrality and Superneutrality of Money

- Two classical concepts describe how monetary policy affects real variables in the long run:
  - **Money neutrality**: a one-time, permanent change in the *level* of the money supply M has no effect on real variables $(c^\star, n^\star, k^\star)$ in steady state. Only nominal variables - prices and wages - change proportionally.
  - **Superneutrality**: a permanent change in the *growth rate* μ of money also has no effect on real variables in steady state. Even sustained inflation leaves real quantities unchanged. Superneutrality is a strictly stronger property than neutrality.
- **In the MIU model**: the three equations governing the real block (c_t, n_t, k_{t+1}) are identical to those of the Classical Growth Model and contain neither M_t nor μ. Hence, the MIU model displays **superneutrality**: both the level and the growth rate of money are irrelevant for real allocations.
- However, superneutrality does **not** imply that monetary policy is irrelevant for welfare. Since real balances $m^\star$ enter utility and depend on μ, the choice of money growth rate affects household welfare even though it leaves $(c^\star, n^\star, k^\star)$ unchanged. This is the **inflation tax**.

The Friedman Rule

- Since higher money growth μ raises $i^\star$, which reduces $m^\star = \chi c^\star / i^\star$, and real balances enter utility, welfare is **decreasing in** μ. The government therefore faces a meaningful policy problem: what is the welfare-maximizing money growth rate?

- Households reach full satiation in liquidity when the opportunity cost of holding money is zero, i.e., when $i^\star = 0$. Using the steady-state Fisher equation:

$$1 + i^\star = \frac{1+\mu}{\beta} = 1 \implies \mu^{FR} = \beta - 1 < 0.$$

- This is the **Friedman Rule** (Friedman, 1969): the optimal monetary policy prescribes a *mild deflation* at rate $\beta - 1$, so that the nominal return on money equals the real return on capital and the distortion from holding money disappears.

- At the Friedman Rule, the money demand condition $m^\star = \chi c^\star / i^\star$ implies $m^\star \to \infty$: households choose to hold unbounded real balances at zero opportunity cost.

- Note: the Friedman Rule achieves the **social optimum** in the MIU model - it eliminates the inflation tax without distorting the real allocation.

Limitations of the MIU Model

- The MIU model is tractable and delivers clean results, but it has serious conceptual and empirical shortcomings.
  - **Circular reasoning**: the model assumes money is useful without explaining why. It takes as a primitive the very thing a theory of money should derive. In particular, it is silent on what makes money different from any other asset that could be placed in the utility function.
  - **Superneutrality is counterfactual**: the model predicts that sustained changes in money growth have no effect on real variables. However, a large empirical literature finds that inflation does affect investment, the capital stock, and labor supply - especially at high inflation rates.
  - **The Friedman Rule is difficult to implement**: the optimal policy prescribes permanent deflation at rate $\beta - 1 \approx -4\%$. In practice, deflation can be destabilizing: it raises the real burden of nominal debt, may cause a deflationary spiral, and runs into the zero lower bound on nominal interest rates. No central bank in the world targets deflation.
  - **Infinite money demand at the optimum**: at the Friedman Rule, $m^\star = \chi c^\star / i^\star \to \infty$. This is an artifact of the log utility specification for money and has no empirical counterpart.

Cash-in-Advance: Overview

- The **Cash-in-Advance (CIA)** model, due to Lucas and Stokey (1983), takes a different approach to introducing money: rather than placing it in the utility function, it imposes a constraint that households must hold money *before* they can purchase consumption goods.

- The CIA constraint captures the idea that trade requires a medium of exchange. You cannot buy groceries with capital or promises - you need cash in hand. This is a more explicit friction than MIU, though still reduced form.

- The model is identical to the Classical Growth Model except for two modifications:
  1. Households face a CIA constraint: real money holdings must cover consumption expenditure.
  2. A government controls the money supply and makes lump-sum transfers, exactly as in the MIU model.

- Crucially, money does **not** appear in the utility function. The demand for money arises entirely from the CIA constraint.

CIA: Households

- Households maximize the same utility as in the Classical Growth Model (no money in utility). They solve:

$$\max_{\{c_t,\, n_t,\, k_{t+1},\, m_{t+1}\}} \sum_{t=0}^{\infty} \beta^t \left[\log(c_t) - \eta \frac{n_t^{1+\nu}}{1+\nu} \right] \quad \text{subject to}$$

$$c_t + k_{t+1} + m_{t+1}(1+\pi_{t+1}) = w_t n_t + R_t k_t + (1-\delta)k_t + m_t + \tau_t, \quad \text{(BC)}$$

$$c_t \leq m_t, \quad \text{(CIA)}$$

$$c_t, \;\; k_{t+1}, \;\; m_{t+1} \geq 0, \qquad k_0, \;\; m_0 \;\; \text{given.}$$

- The CIA constraint says that consumption in period t cannot exceed the real money balances m_t the household carries into that period. If $c_t < m_t$, excess money is held idle - a dominated strategy when $i > 0$, so the constraint will bind.

CIA: Households

- The Lagrangian attaches multiplier λ_t to the budget constraint and $\mu_t \geq 0$ to the CIA constraint:

$$\begin{aligned}\mathcal{L} = &\sum_{t=0}^{\infty} \beta^t \left[\log(c_t) - \eta \frac{n_t^{1+\nu}}{1+\nu} \right] \\ &+ \sum_{t=0}^{\infty} \lambda_t \left[w_t n_t + R_t k_t + (1-\delta) k_t + m_t + \tau_t - c_t - k_{t+1} - m_{t+1}(1+\pi_{t+1}) \right] \\ &+ \sum_{t=0}^{\infty} \mu_t \left[m_t - c_t \right]\end{aligned}$$

- The multiplier λ_t is the shadow value of wealth: one extra unit of income can be used to buy capital, money, or more consumption. The multiplier μ_t is the **additional** value of having money beyond its financial value - it relaxes the CIA constraint, allowing more consumption.

- By complementary slackness: $\mu_t \geq 0$ and $\mu_t(m_t - c_t) = 0$. So either the CIA constraint binds ($c_t = m_t$, $\mu_t > 0$) or money is held in excess ($c_t < m_t$, $\mu_t = 0$).

CIA: Households

- The FOCs of the household are:

$$\frac{\partial \mathcal{L}}{\partial c_t} = 0 \Rightarrow \frac{\beta^t}{c_t} = \lambda_t + \mu_t,$$
$$\frac{\partial \mathcal{L}}{\partial n_t} = 0 \Rightarrow \eta \beta^t n_t^{\nu} = w_t \lambda_t,$$
$$\frac{\partial \mathcal{L}}{\partial k_{t+1}} = 0 \Rightarrow \lambda_t = \left[R_{t+1} + 1 - \delta\right] \lambda_{t+1},$$
$$\frac{\partial \mathcal{L}}{\partial m_{t+1}} = 0 \Rightarrow \lambda_t (1 + \pi_{t+1}) = \lambda_{t+1} + \mu_{t+1}.$$

- Compare these with the MIU model. The Euler for capital (third condition) is **identical**. The key difference is in the consumption FOC: the total cost of consuming includes both λ_t (budget cost) and μ_t (tightening the CIA constraint). The fourth condition now says the cost of acquiring money equals its financial value next period plus its CIA value next period.

CIA: Deriving the Euler Equation (1/2) - The Shadow Value of Money

- Goal: combine the four FOCs into a single Euler equation in consumption, eliminating the CIA multiplier μ_t along the way.
- The consumption FOC gives $\mu_t = \beta^t/c_t - \lambda_t$; one period ahead, $\mu_{t+1} = \beta^{t+1}/c_{t+1} - \lambda_{t+1}$.
- Substitute into the money FOC, $\lambda_t(1+\pi_{t+1}) = \lambda_{t+1} + \mu_{t+1}$: the λ_{t+1} terms cancel, leaving a clean expression for λ_t:

$$\lambda_t(1+\pi_{t+1}) = \lambda_{t+1} + \left[\frac{\beta^{t+1}}{c_{t+1}} - \lambda_{t+1}\right] \implies \lambda_t = \frac{\beta^{t+1}}{(1+\pi_{t+1})c_{t+1}}. \quad (\star)$$

- Equation $(\star)$ says: the shadow value of wealth today equals tomorrow's marginal utility of consumption, deflated by tomorrow's inflation. Since it holds for every t, it also holds one period later:

$$\lambda_{t+1} = \frac{\beta^{t+2}}{(1+\pi_{t+2})c_{t+2}}.$$

CIA: Deriving the Euler Equation (2/2) - Combining with the Capital FOC

- Plug the two expressions for λ_t and λ_{t+1} from the previous slide into the capital FOC, $\lambda_t = [R_{t+1} + 1 - \delta]\lambda_{t+1}$:

$$\frac{\beta^{t+1}}{(1+\pi_{t+1})c_{t+1}} = [R_{t+1} + 1 - \delta] \frac{\beta^{t+2}}{(1+\pi_{t+2})c_{t+2}}.$$

- Cross-multiply and cancel one power of β to solve for c_{t+2}/c_{t+1}:

$$\frac{c_{t+2}}{\beta\, c_{t+1}} = [R_{t+1} + 1 - \delta] \frac{1+\pi_{t+1}}{1+\pi_{t+2}}.$$

- This holds for every t. Relabeling the time index ($t \to t-1$) turns it into a relation between c_{t+1} and c_t - the **Euler equation**:

$$\boxed{\frac{c_{t+1}}{\beta\, c_t} = [R_t + 1 - \delta] \frac{1+\pi_t}{1+\pi_{t+1}}.}$$

CIA: Intuition - Why Does Inflation Enter the Euler Equation?

- In the Classical Growth Model and in MIU, the capital Euler equation is purely **real**: $c_{t+1}/(\beta c_t) = R_{t+1} + 1 - \delta$. Money is irrelevant to how households trade off consumption today versus tomorrow.
- In CIA, the Euler equation carries an inflation term, $\dfrac{1+\pi_t}{1+\pi_{t+1}}$, because $c_t \leq m_t$ binds: consuming *requires* holding money one period in advance. Wealth saved today becomes tomorrow's consumption only by first passing through money, whose real purchasing power is eroded by inflation while held. This is what $(\star)$ showed: λ_t is tied to *next period's* marginal utility of consumption, discounted for the inflation that erodes the money used to buy it.
- In MIU, money enters utility **directly** with its own FOC - it never intermediates the purchase of goods, so the consumption FOC has no μ_t term, and the capital Euler equation is untouched by inflation.
- Subtlety: in a **steady state**, $\pi_t = \pi_{t+1} = \mu$, so $(1+\pi_t)/(1+\pi_{t+1}) = 1$ and the Euler condition reverts to $R^\star + 1 - \delta = 1/\beta$ - same as MIU. Superneutrality fails in CIA not through this condition, but through the **intratemporal** one, where μ enters via $i^\star$ (next slide).

CIA: Derived Conditions

- Combining the four FOCs (see the previous two slides for the step-by-step derivation) gives the **Euler equation**:

$$\frac{c_{t+1}}{\beta c_t} = (R_t + 1 - \delta)\frac{1+\pi_t}{1+\pi_{t+1}}.$$

  In steady state ($\pi_t = \pi_{t+1} = \mu$) this simplifies to the same condition as before: $R^\star + 1 - \delta = 1/\beta$.

- Combining the consumption and labor FOCs, and using $\mu_t/\lambda_t = i_{t+1}$ in steady state (please verify), the **intratemporal condition** becomes:

$$\eta c^\star (n^\star)^\nu (1 + i^\star) = w^\star.$$

- The nominal interest rate acts as a **tax on consumption**: it raises the effective price of consuming relative to leisure, since buying consumption goods requires holding money with opportunity cost $i^\star$.

- Finally, when the CIA constraint binds, **money demand is simply**:

$$m^\star = c^\star.$$

CIA: Steady State

- In a steady state, all variables are constant: $c_t = c^\star$, $n_t = n^\star$, $k_t = k^\star$, $m_t = m^\star$, $\pi_t = \pi^\star = \mu$. The steady-state system is (make sure you understand where each equation comes from):

$$\eta c^\star (n^\star)^\nu (1 + i^\star) = (1 - \alpha) A \left[\frac{k^\star}{n^\star}\right]^\alpha,$$

$$\frac{1}{\beta} = \alpha A \left[\frac{n^\star}{k^\star}\right]^{1-\alpha} + 1 - \delta,$$

$$c^\star = A(k^\star)^\alpha (n^\star)^{1-\alpha} - \delta k^\star,$$

$$m^\star = c^\star, \qquad 1 + i^\star = \frac{1 + \mu}{\beta}.$$

- The last equation links the nominal interest rate $i^\star$ to the money growth rate μ set by the government. Notice that $i^\star$ then enters the first equation - the intratemporal condition. This is the channel through which monetary policy affects the real economy.

CIA: Failure of Superneutrality

- In the Classical Growth Model and in the MIU model, the steady-state real allocation $(c^\star, n^\star, k^\star)$ is determined by the Euler equation, the intratemporal condition, and the resource constraint - none of which depend on μ.

- In the CIA model, μ enters the intratemporal condition through $i^\star = (1+\mu)/\beta - 1$. A higher money growth rate therefore **changes the real allocation**. The CIA model **fails superneutrality**.

- The economic mechanism is straightforward: higher μ raises $i^\star$, which increases the implicit tax on consumption (since buying goods requires money with opportunity cost $i^\star$). Households respond by substituting away from consumption toward leisure, which reduces labor supply $n^\star$, capital accumulation $k^\star$, and output $Y^\star$.

- The **Friedman Rule** ($i^\star = 0$, i.e., $\mu = \beta - 1$) eliminates this distortion entirely, restoring the same real allocation as the Classical Growth Model. In CIA, the Friedman Rule is a **real efficiency** result - not merely a welfare improvement as in MIU.

CIA vs. MIU: A Comparison

- Both models deliver a demand for money and a role for monetary policy, but they differ in important ways:
  - **Source of money demand**: in MIU, money is demanded because it yields utility. In CIA, it is demanded because it is required to buy goods. CIA offers a more explicit friction, but both are still reduced-form.
  - **Superneutrality**: MIU satisfies superneutrality - inflation has no real effects. CIA does not - higher inflation acts as a consumption tax, reducing labor, capital, and output. CIA is more consistent with empirical evidence.
  - **Money demand**: in MIU, $m^\star = \chi c^\star / i^\star$ decreases with the nominal interest rate. In CIA, $m^\star = c^\star$ is independent of $i^\star$ (as long as the CIA binds).
  - **Friedman Rule**: both models recommend $\mu = \beta - 1$. In MIU, this is a welfare improvement with no real effects. In CIA, it also restores the first-best real allocation.

OLG Models

- Up to this moment, we have considered models in which there is a mass of identical consumers (same preferences and endowments). We focused our attention on an infinitely lived representative agent.
  - Both models are useful to explain some features of the data.
  - But in the context of these models can we analyze central policy topics like inequality, life cycle, inefficiencies?
  - NO! Why?
- The Overlapping Generations Model (OLG) is a mayor workhorse in the macro literature that will provide a natural context to analyze issues that the classical model cannot explain.
  - Pioneered by Allais (1947), Samuelson (1958), and Diamond (1965).
  - Appealing since it allows us to integrate micro and macro data.
  - We will be able to analyze topics like: social security, source of market inefficiencies, the value of money, distributive effect of taxes, life cycle savings, among others.

Basic Ingredients of an OLG

- Time is discrete $t = 1, 2, ...$
- We will consider a single nonstorable consumption good.
- **Key feature**: individuals do not live forever.
  - For the moment, we will assume that individuals only live for two periods and then die.
  - Every period a new generation is born, with mass 1.
  - We will index generations by the period they are born: g^t will be the generation that was born at period t.
  - Each generation will be endowed with units of consumption every period it lives. Let (e_t^t, e_{t+1}^t) be the endowment that generation t receives at period t and at $t+1$.

Consumption Timing

- We denote (c_t^t, c_{t+1}^t) the consumption of generation t at periods t and $t+1$.
- Notice that every period t there are two generations alive:
  - **Young Generation**: born at period t, endowed with e_t^t who consume c_t^t.
  - **Old Generation**: born at period $t-1$, endowed with e_t^{t-1} who consume c_t^{t-1}.
- There is an **Initial Old Generation** at period $t=1$ who is endowed e_1^0 and consumes c_1^0.

Consumption Timing

generation\time	1	2	...	t	$t+1$
0	(c_1^0, e_1^0)				
1	(c_1^1, e_1^1)	(c_2^1, e_2^1)			
$\vdots$			$\ddots$		
$t-1$				(c_t^{t-1}, e_t^{t-1})	
t				(c_t^t, e_t^t)	(c_{t+1}^t, e_{t+1}^t)
$t+1$					$(c_{t+1}^{t+1}, e_{t+1}^{t+1})$

Figure: Timing in an OLG Model

- Alternative and interesting interpretation: an OLG model can be seen as a framework where there are infinitely lived consumers that **only care about consumption in two periods**.
  - By reading the model in this way, we can see that in an OLG context there are **Incomplete Markets**!
  - This suggests that Pareto Optimality may not hold in OLG models.
  - We will see that, in many cases, equilibria are inefficient.

Preferences and Feasibility

- Preferences of generation t are given by:

$$U_t(c) = log(c_t^t) + \beta log(c_{t+1}^t).$$

- The initial old have preferences $U_0(c) = log(c_1^0)$.

- We say a consumption allocation $c_1^0, \{(c_t^t, c_{t+1}^t)\}_{t=1}^{\infty}$ is **feasible** if:

$$c_t^{t-1} + c_t^t = e_t^{t-1} + e_t^t \quad \text{for every} \quad t \geq 1.$$

Selfishness

- What are the incentives for each generation?
  - The young generation would like to consume (obviously), however, it knows that they will live for one more period. Hence, young people have incentives to save for the future.
  - The old generation also values consumption and they die the next period, hence they **do not have incentives to do anything but consume**.
  - Up until this point, even though the young generation will like to save for the future they cannot do it since there is no one to trade with!
  - If we leave the model like this it will always imply an **autarkic equilibrium**, where generations only consume their endowments and there is no trade.
- If we want to induce trade between generations (which is a feature we would like since trade usually increases welfare) we need to introduce a mechanism that promotes it.
  - Solution: **Money**!!

Money and OLG

- Why money promotes trade?
  - Simple: it is a **promise of future payment**.
  - If the initial old have money to trade, they can give it to the initial young at $t = 1$, which they when old ($t = 2$) can trade with the young in exchange for consumption goods.
  - Hence, money is a certificate that guarantees you as a young person that, when old, you will receive consumption goods in exchange for it.
  - Money is memory (Kocherlakota, 1998).
- In this model, money is a **bubble**: even though it has no intrinsic value (it is literally a piece of paper), people use it to trade and the market valuates it at a positive price.
  - This can only occur in the context of incomplete markets. There is a famous theorem by Samuelson which states that in an economy with infinitely lived agents and no frictions money is never valued in equilibrium.
- What is the role of monetary policy then?

Market Structure

- We will mainly focus on markets that allow sequential trading. Hence, there are markets for the consumption good in every period.
- In addition, there is an asset through which generations can do their savings.
- We will denote a_t^t the savings done by generation t which will be delivered/payed in period $t+1$.
- Let r_{t+1} be the interest rate related to such assets between period t and $t+1$.

Market Equilibrium

- Let $m \geq 0$ be the amount of money available in the economy.
- An equilibrium is an allocation $\hat{c}_1^0, \{(\hat{c}_t^t, \hat{c}_{t+1}^t, \hat{a}_t^t)\}_{t=1}^{\infty}$ and a sequence of interest rates $\{\hat{r}_t\}_{t=1}^{\infty}$ such that:
  1. Given interest rates, $(\hat{c}_t^t, \hat{c}_{t+1}^t, \hat{a}_t^t)$ solves generation's t problem:

$$\max_{\{c_t^t, c_{t+1}^t, a_t^t\}} \quad log(c_t^t) + \beta log(c_{t+1}^t) \quad \text{subject to}$$
$$c_t^t + a_t^t \leq e_t^t,$$
$$c_{t+1}^t \leq e_{t+1}^t + (1 + \hat{r}_{t+1}) a_t^t.$$

  2. Given interest rates, $\hat{c}_1^0$ solves the initial old problem:

$$\max_{\{c_1^0\}} \quad log(c_1^0) \quad \text{subject to}$$
$$c_1^0 \leq e_1^0 + (1 + \hat{r}_1) m.$$

  3. For all $t \geq 1$ markets clear:

$$\hat{c}_t^{t-1} + \hat{c}_t^t = e_t^{t-1} + e_t^t.$$

What are Savings in this Economy?

- At every period:
$$c_t^t + a_t^t = e_t^t \quad c_t^{t-1} = e_t^{t-1} + (1 + r_t)a_{t-1}^{t-1}.$$

- If we add these budget constraints:
$$c_t^{t-1} + c_t^t + a_t^t = e_t^{t-1} + e_t^t + (1 + r_t)a_{t-1}^{t-1},$$
which then by market clearing becomes:
$$a_t^t = (1 + r_t)a_{t-1}^{t-1}.$$

- In particular, at $t = 1$:
$$a_1^1 = (1 + r_1)m,$$
which implies that:
$$a_t^t = \prod_{\tau=1}^{t}(1 + r_\tau)m.$$

- Then r measures the **return of money**.

Computing Prices

- Remember that p_t is the price of consumption in period t in terms of the initial period goods.
- How can we recover these prices?
  - Easy! As we have been doing all this time:

$$\frac{p_t}{p_{t+1}} = 1 + r_{t+1} \quad \frac{1}{p_1} = 1 + r_1$$

- Once we have these prices, we know that the consumption allocation is the same in either market structure.
- We can give an additional interpretation to this relationship between p and r in this model:
  - Remember r_{t+1} is the return on money.
  - In this case:

$$1 + r_{t+1} = \frac{p_t}{p_{t+1}} = \frac{1}{1 + \pi_{t+1}}.$$

  - Then (using a log approximation) $r_{t+1} = -\pi_{t+1}$.

Offer Curves

- How do we actually find an equilibrium for this economy?
- Gale (1973) proposed a nice algorithm to do this.
- We will assume that the endowment for the young generations is always e_1 while for the old it is always e_2.
- It will be convenient to work with the price ratio $P_t \equiv \dfrac{p_t}{p_{t+1}}$.
- Let $c_t^t(P_t), c_{t+1}^t(P_t)$ denote the optimal consumption of generation t at periods t and $t+1$. **Recall** that these demands are *given* by:

$$c_t^t(P_t) = \frac{e_1 + e_2/P_t}{1+\beta}, \quad c_{t+1}^t(P_t) = \frac{\beta(e_1 P_t + e_2)}{1+\beta}.$$

- Let us consider the (Walrasian) excess demand functions:

$$y^t(P_t) = c_t^t(P_t) - e_1, \quad z^t(P_t) = c_{t+1}^t(P_t) - e_2.$$

Offer Curves

- As we vary p_{t+1}/p_t we obtain what Gale called **Offer Curve**.

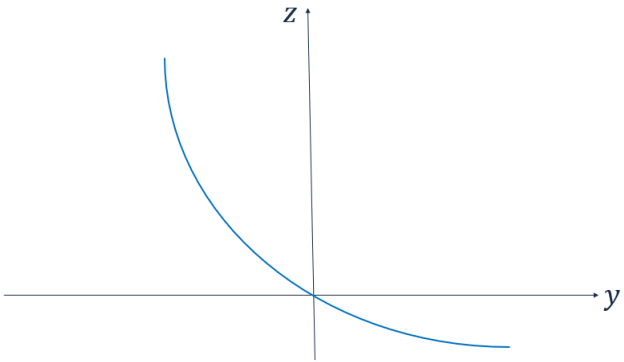


Figure: Offer Curve

Offer Curves

- The excess demand functions must satisfy the budget constraint:

$$p_t y^t(P_t) + p_{t+1} z^t(P_t) = 0 \implies \frac{z^t(P_t)}{y^t(P_t)} = -P_t.$$

- For the initial old:

$$z^0(p_1) = \frac{m}{p_1}.$$

- Then, given a choice of p_1 we know z^0, and the offer curve lets us iterate to get the entire future sequence of prices and allocations!

Deriving the Offer Curve: General Formula

- Let's derive the offer curve **once**, in general, for $\beta = 1$ (the case we will use throughout). From the given demands:

$$y = \frac{e_1 + e_2/P}{2} - e_1 = \frac{e_2}{2P} - \frac{e_1}{2}, \quad z = \frac{e_1 P + e_2}{2} - e_2 = \frac{e_1 P}{2} - \frac{e_2}{2}.$$

- **Step 1**: solve the y-equation for P:

$$P = \frac{e_2}{2y + e_1}.$$

- **Step 2**: substitute into the z-equation and simplify:

$$z = \frac{e_1}{2}\left(\frac{e_2}{2y + e_1}\right) - \frac{e_2}{2} = \frac{e_1 e_2 - e_2(2y + e_1)}{2(2y + e_1)} = -\frac{e_2\, y}{2y + e_1}.$$

- **General Offer Curve Formula** (for $\beta = 1$):

$$\boxed{z = -\frac{e_2\, y}{2y + e_1}}$$

Offer Curves: Computing Equilibrium Prices

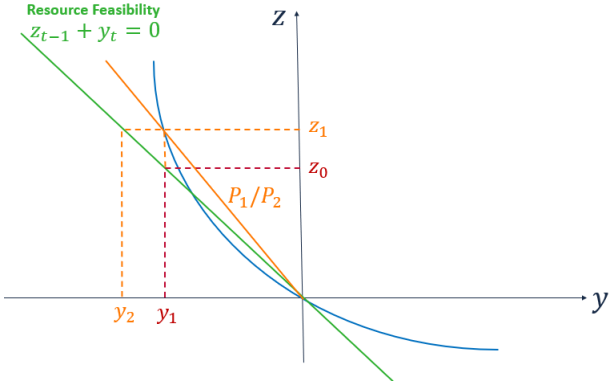


Figure: Offer Curve

Offer Curves: Some Important Remarks

- Notice that if we change m the entire sequence of prices and allocations is shifted!
- Then, we can see that in an OLG model there is a multiplicity of equilibria.
  - This is not a desirable property of a model.
  - Which of the (possibly infinite number) of equilibria is "the best"?
  - Are all equilibria efficient? Seems unlikely.
- What is the role of monetary policy?

The Balasko-Shell Theorem

- When is an equilibrium efficient?
- Let $\{\hat{p}_t\}_{t=1}^{\infty}$ denote the sequence of equilibrium prices.
  - The **Balasko-Shell Theorem** states that $\{\hat{p}_t\}_{t=1}^{\infty}$ induces an equilibrium allocation that is Pareto Efficient if and only if
$$\sum_{t=1}^{\infty} p_t < \infty.$$
- Hence, if prices "explode", $\sum_{t=1}^{\infty} p_t = \infty$, the implied equilibrium is not efficient.
  - Remember that prices and interest rates (return of money) are inversely related in equilibrium.
  - Hence, prices exploding in equilibrium is equivalent to **interest rates being too low (or even negative)**.
  - **Key Intuition**: If interest rates are too low, trade is desincentivized and hence, the economy is "close" to the autarkic equilibrium, that we know is inefficient (in general).

Numerical Example

- Imagine an OLG economy where $e_t^t = 30$ and $e_{t+1}^t = 4$ for every t. Preferences are given by $log(c_t^t) + log(c_{t+1}^t)$.

  1. Set up the problem that the consumer solves, and state (recall, don't re-derive) the consumption demands.
  2. Suppose there is no money in this economy. What are the equilibrium prices and allocations?
  3. Is this equilibrium efficient?
  4. Using the offer curve recipe, find the Offer Curve equation for this economy.
  5. Suppose $m = 10$ and $p_1 = 1$. Find the sequence of equilibrium prices.
  6. Is this equilibrium Pareto Optimal?

Numerical Example: Question 1

- Generation t solves:

$$\max_{c_t^t, c_{t+1}^t} \; log(c_t^t) + log(c_{t+1}^t) \;\text{ subject to}$$

$$p_t c_t^t + p_{t+1} c_{t+1}^t = 30 p_t + 4 p_{t+1}.$$

- Here $\beta = 1$, $e_1 = 30$, $e_2 = 4$. We are **given** the demands (recall the general formula from the previous section – no need to re-derive):

$$c_t^t(P_t) = \frac{30 + 4/P_t}{2} = 15 + \frac{2}{P_t}, \quad c_{t+1}^t(P_t) = \frac{30P_t + 4}{2} = 15P_t + 2,$$

where $P_t = p_t / p_{t+1}$.

Numerical Example: Question 2

- Remember, if $m = 0$ then there is no trade in this economy (why?). Then, in equilibrium each generation will consume its endowment. Hence:

$$c_t^t = 30 \quad c_{t+1}^t = 4.$$

- Using any of the demand functions, we can attain the equilibrium price ratio:

$$30 = c_t^t = 15 + \frac{2}{P_t},$$

implying that:

$$P_t = \frac{p_t}{p_{t+1}} = \frac{2}{15} \implies \frac{p_{t+1}}{p_t} = \frac{15}{2},$$

and hence:

$$p_t = \left[\frac{15}{2}\right]^t.$$

Numerical Example: Question 3

- To assert if this equilibrium is or not Pareto efficient, we need to compute:

$$\sum_{t=1}^{\infty} p_t = \sum_{t=1}^{\infty} \left[\frac{15}{2}\right]^t = \infty,$$

since $\frac{15}{2} > 1$.

- Hence, by the Balasko-Shell Theorem, this equilibrium is not efficient.
- Intuition?

Numerical Example: Question 4

- We already derived the general offer curve formula (for $\beta = 1$):

$$z = -\frac{e_2\, y}{2y + e_1}.$$

- Plugging in $e_1 = 30$, $e_2 = 4$:

$$z = -\frac{4y}{2y + 30} = -\frac{2y}{y + 15}.$$

- No need to redo the P-substitution algebra each time – just plug e_1, e_2 into the recipe.

Numerical Example: Question 5

- Now, since $m = 10, p_1 = 1$ then $z^0 = 10$. By market clearing conditions this implies that $y^1 = -10$.
- Then, by the offer curve and the price equation:

$$z^1 = \frac{2(10)}{-10+15} = 4 \quad P_1 = \frac{p_1}{p_2} = \frac{1}{p_2} = \frac{-z^1}{y^1} = \frac{2}{5}.$$

- Now, since $z^1 = 4$ this implies that $y^2 = -4$ and hence:

$$z^2 = \frac{2(4)}{-4+15} = \frac{8}{11} \quad P_2 = \frac{p_2}{p_3} = \frac{-z^2}{y^2} = \frac{2}{11}.$$

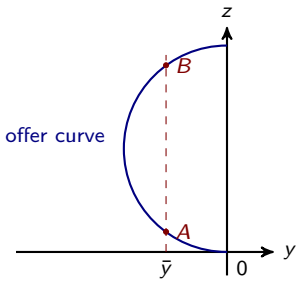
- This gives a simple recipe to iterate forward at any t: given z^{t-1}, set $y^t = -z^{t-1}$, then plug into the offer curve to get $z^t = \frac{2(-y^t)}{y^t+15}$ and $P_t = -z^t/y^t$.
- Repeating this, both $|y^t|$ and $|z^t|$ shrink toward 0 while the denominators grow – the economy is converging back to autarky. (We will use this pattern, not the exact fractions, to discuss efficiency.)

Multiplicity, Take Two: Even for Fixed m and p_1

- We already saw that different choices of m support different equilibria. But the problem can run deeper than that.
- In our numerical example, the offer curve was a **monotonic**, single-valued function of y: each y^t pinned down exactly one z^t. That is precisely why, once m and p_1 were fixed, Gale's algorithm gave a *unique* future price path.
- But the offer curve is really the locus traced out by the price ratio P: $\{(y(P), z(P)) : P > 0\}$. Nothing guarantees $y(P)$ is monotonic in P – for more general (still perfectly standard, well-behaved) preferences, it need not be.
- If $y(P)$ folds back on itself, a single value of y^t can correspond to **more than one** point on the offer curve – more than one z^t, more than one P_t, more than one p_{t+1}.
- Each such point launches its own self-consistent future path: genuinely different equilibria, for *identical* fundamentals, *identical* m, and *identical* p_1.

A Backward-Bending Offer Curve

- Here is what that looks like: an offer curve that folds back on itself instead of moving monotonically.



- At $y^1 = \bar{y}$ – pinned down by our fixed m and p_1 – **both** A and B lie on the offer curve.
- $z_A \neq z_B$ means a different P_1, a different p_2, and an entirely different future path from that point on.
- Both are legitimate equilibria: nothing in the fundamentals tells us which one is "the" equilibrium.

Sunspots and the Coordination Problem

- If more than one equilibrium is consistent with the exact same fundamentals, which one actually happens?
- The fundamentals $(e_1, e_2, \beta, m, p_1)$ do not answer this – **expectations** do. Whatever agents believe will happen tends to make itself true.
- This opens the door to **sunspot equilibria** (Azariadis, 1981; Cass and Shell, 1983): agents condition their behavior on some extrinsic, payoff-irrelevant signal (e.g., "if the sunspot says Heads, play equilibrium A; if Tails, play B"). Since each branch remains self-fulfilling, this randomization is *itself* a valid rational-expectations equilibrium.
  - Real volatility in prices and allocations, with **zero** change in fundamentals!
- The literature has proposed various "fixes": stability/learnability refinements (select the equilibrium that adaptive learning would converge to), or policy rules designed to pin down a unique equilibrium.
- At its core, this is a **coordination problem**: agents must somehow agree on which equilibrium to play, even though nothing in the model forces that agreement. We will return to this question – and to how monetary policy design can help resolve it – later in the course.

Discussing the OLG Model Assumptions

- As we have seen, the particular assumptions made for the OLG framework have some interesting consequences.
- How important are these assumptions in order for the results to still hold?
- In particular, we will discuss the role of:
  1. Households being selfish vs being altruistic.
  2. Agents living a finite number of periods vs living forever.
  3. Agents knowing in which period they die vs not knowing this.

Altruism

- One could criticize the OLG model since many of its interesting results rely on the "selfishness" feature of the model.

- What happens if we allow room for altruism?

  - In particular, we could think that generation $t-1$ are the parents of generation t, and hence generation $t-1$ cares on the well being of its children.
  - We introduce the notion of **bequest**: units of consumption good a parent (old generation $t-1$) can give its children (young generation t) after they die.
  - We assume the bequest is non-strategic.
  - Let $b_t^{t-1} \geq 0$ be the bequest generation $t-1$ leaves to generation t. Then, in period $t+1$ generation t (which is now old) will have as income:

$$e_{t+1}^t(1+r_{t+1})a_{t+1}^t + b_t^{t-1}.$$

Altruism

- Then, parents solve the following problem (notice that their parents, generation $t-2$, left them a bequest of size b_t^{t-1} which they take as **given**):

$$V_{t-1}(b_t^{t-1}) = \max_{\{c_{t-1}^{t-1}, c_t^{t-1}, b_{t+1}^t, a_t^{t-1}\}} \left\{ log(c_{t-1}^{t-1}) + \beta log(c_t^{t-1}) + \alpha V_t(b_{t+1}^t) \right\}$$

subject to

$$c_{t-1}^{t-1} + a_t^{t-1} = e_{t-1}^{t-1},$$

$$c_t^{t-1} + b_{t+1}^t = e_t^{t-1} + (1 + r_t) a_t^{t-1} + b_t^{t-1}$$

- Where we can interpret $0 < \alpha < 1$ as the weight parents give in their own utility to the well-being of their children.

Altruism

- But then, starting from the initial old, notice the following (I will not write the budget restrictions, but remember they are there):

$$V_0(b_1) = \max\left\{log(c_1^0) + \alpha V_1(b_2^1)\right\}$$

$$= \max\left\{log(c_1^0) + \alpha \max\left\{log(c_1^1) + \beta log(c_2^1) + \alpha V_2(b_3^2)\right\}\right\}$$

$$= \max\left\{log(c_1^0) + \alpha\left[log(c_1^1) + \beta log(c_2^1)\right] + \alpha^2 V_2(b_3^2)\right\}$$

$$= \max\left\{log(c_1^0) + \alpha\left[log(c_1^1) + \beta log(c_2^1)\right] + \alpha^2 \max\left\{log(c_2^2) + \beta log(c_3^2) + \alpha V_3(b_4^3)\right\}\right\}$$

$$= \max\left\{log(c_1^0) + \alpha\left[log(c_1^1) + \beta log(c_2^1)\right] + \alpha^2\left[log(c_2^2) + \beta log(c_3^2)\right] + \alpha^3 V_4(b_5^4)\right\} ...$$

$$= \max\left\{\sum_{t=0}^{\infty} \alpha^t \tilde{u}(c_t)\right\}$$

- Hence, allowing altruism when parents care $\alpha > 0$ about their children's utility leads to the same problem that an infinitely lived representative agent solves! WOW!

Altruism: Final Remarks

- Then, assuming altruism leads households to behave **AS IF** they were infinitely lived, even though they actually live for two periods.

- This result is another justification commonly used in the literature for the infinitely lived representative agent assumption.

- Shall we use selfish or altruistic households?
  - **It depends on what you wish to highlight**!!
  - Assuming altruism is equivalent of assuming a representative agent, hence you recover Pareto Optimality of equilibria, complete markets, etc...
  - However, assuming altruistic households takes you away from an OLG framework. Then you can no longer rationalize inequality, the value of money, monetary policy, incomplete markets, etc...

- Important lesson: one has to seriously think about the assumptions one is making, since as you can now understand, they have important consequences!

Finite vs Infinite Life

- Another key assumption in the OLG framework is that **agents do not live forever**.
- Assuming this creates an heterogeneity within each period, since there are agents who are at different moments of their life cycle, and therefore have different incentives.
  - For example, in a world in which agents live for two periods, young agents have an incentive to save while old agents have an incentive to consume.
  - This creates a tension within the model, leading to incomplete markets, and therefore makes interesting to introduce mechanisms such as money.
- On the other hand, if agents lived forever, even in a context in which we have heterogeneous agents (e.g. a continuum of agents each with different endowments), they would have incentives to trade either with their "future selves" or with other agents.
  - This takes us back to the world of complete markets.

Outline

1. A Primer on Monetary Economics
2. Foundations of Money
3. The Welfare Cost of Inflation
4. Deeper Foundations of Money
5. Quantity Theory, Expectations, and Policy
6. Time Inconsistency and Central Bank Design
7. Monetary-Fiscal Interactions
8. The Role of Private Banks

Motivation

- Price stability is the primary mandate of most central banks in the world. But is inflation actually costly? If so, *how* costly, and through *what channels*?

- The three models we have studied suggest three qualitatively different answers:
  - **MIU**: inflation erodes real money holdings, destroying liquidity services that households value. The cost is a pure consumption-equivalent loss with *no real distortion*.
  - **CIA**: inflation acts as a tax on consumption, distorting the labor-leisure decision and reducing capital accumulation. The cost includes a *real output loss*.
  - **OLG**: inflation reduces the real return on money, undermining its role as a store of value and distorting *intergenerational trade*. The cost may include the collapse of the monetary equilibrium itself.

- The canonical framework for measuring this cost is the **Bailey triangle** (Bailey, 1956). We study it carefully, and then discuss what each model adds beyond it.

Inflation as a Tax: The Bailey Insight

- Martin Bailey (1956) made a deceptively simple observation: **inflation is an excise tax on money balances**.

- By holding money rather than capital, households forgo the real return R. With inflation π, money loses purchasing power at rate π. The total opportunity cost of holding one unit of real balances for one period is the **nominal interest rate**:

$$i \approx R + \pi.$$

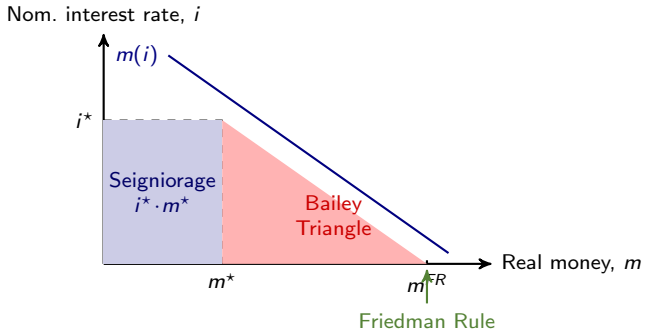
- Households respond rationally: they economize on cash. Just as a tax on any good reduces its consumption below the social optimum, a higher nominal interest rate reduces money holdings below the level that would prevail under the Friedman Rule ($i = 0$).

- This reduction in money holdings is the source of the welfare loss. Households incur real costs - more frequent bank trips, foregone transactions - to substitute away from money. These costs constitute the **deadweight loss** of the inflation tax.

- Crucially, this deadweight loss can be read directly off the money demand curve, just as the DWL of an excise tax is read off the demand curve for the taxed good.

The Bailey Triangle

- Consider the money demand curve $m(i)$, which is downward-sloping. At the Friedman Rule ($i = 0$), households hold m^{FR}; at the prevailing rate $i^\star$, they hold $m^\star < m^{FR}$.



- **Seigniorage** (blue): the government's revenue from the inflation tax.
- **Bailey Triangle** (red): the pure deadweight loss from under-provision of liquidity. This is the **welfare cost of inflation**.

Empirical Magnitudes: Lucas (2000)

- Lucas (2000, "Inflation and Welfare") provided the most careful empirical measurement of the Bailey triangle for the United States, using annual data on M1/GDP and short-term interest rates from 1900–1994.

- He estimated two money demand specifications and computed the area under each:

	Log-log	**Semi-log**
Welfare cost of $\pi = 10\%$ vs. $\pi = 0\%$	$\approx 0.9\%$ of GDP	$\approx 0.45\%$ of GDP
Welfare cost of $\pi = 0\%$ vs. Friedman Rule	$\approx 0.4\%$ of GDP	$\approx 0.25\%$ of GDP

- These numbers are **small but not negligible**. For the US in 2024, 1% of GDP corresponds to roughly \$280 billion - a permanent annual loss in consumer welfare.

- The estimates are larger for countries with historically high inflation. For $\pi = 100\%$ (not uncommon in Latin America in the 1980s), the welfare cost can reach 3–5% of GDP.

- Important caveat: these are **steady-state** estimates. They do not capture the *transition costs* of disinflation - the recessions associated with tightening monetary policy (the sacrifice ratio).

The Bailey Triangle in the MIU Model

- In the MIU model, the money demand function is $m(i) = \chi c^\star / i$ (derived from the optimal money demand condition). This is highly elastic for low i and inelastic for high i.

- The welfare cost of inflation at rate $i^\star$ relative to the Friedman Rule is the area under the money demand curve above the seigniorage rectangle:

$$\mathcal{W}_{MIU} = \int_0^{i^\star} m(i)\, di - i^\star m^\star = \chi c^\star \ln\left(\frac{i^\star}{\varepsilon}\right) - i^\star m^\star,$$

  where $\varepsilon \to 0$ reflects that money demand diverges as $i \to 0$ under log utility - an artifact we discussed earlier.

- Since superneutrality holds in MIU, the Bailey triangle is the **complete** welfare cost: inflation has no effect on $(c^\star, n^\star, k^\star)$, only on $m^\star$. The triangle misses nothing.

- The welfare cost grows more than proportionally with $i^\star$: doubling the nominal interest rate more than doubles the triangle, because the demand curve is convex. This means going from 2% to 4% inflation is *more* harmful than going from 0% to 2%.

The CIA Model: Why the Triangle is Not Enough

- In the CIA model, money demand is $m^\star = c^\star$ - perfectly tied to consumption, not to the nominal interest rate. There is no portfolio substitution away from money as $i^\star$ rises: households cannot economize on cash relative to consumption, because they *need* cash to buy consumption.

- Naively applying the Bailey triangle to CIA would suggest a very small welfare cost: since $m(i)$ is not directly sensitive to i, the area under the demand curve appears small. But this reasoning is **deeply misleading**.

- The welfare cost of inflation in CIA is **larger** than in MIU, not smaller. The reason is that CIA breaks superneutrality: higher inflation distorts real variables that MIU leaves untouched. The intratemporal condition:

$$\eta\, c^\star (n^\star)^\nu (1 + i^\star) = w^\star$$

  reveals that the nominal interest rate acts as a **tax on consumption relative to leisure**.

- The Bailey triangle therefore captures only part of the story in CIA. The money demand channel is still present, but it is *not the dominant one*.

The CIA Model: The Full Welfare Cost

- The general-equilibrium mechanism in CIA generates a chain of real losses that go well beyond the Bailey triangle. Higher $i^\star$ induces:

$$i^\star \uparrow \implies n^\star \downarrow \implies k^\star \downarrow \implies Y^\star \downarrow \implies c^\star \downarrow .$$

- The full welfare cost in CIA therefore has **two distinct components**:
  1. The **Bailey triangle**: the welfare loss from holding fewer real balances than at the Friedman Rule. This is the same channel as in MIU, and can be read off the money demand curve.
  2. A **general-equilibrium distortion**: the welfare loss from the contraction of labor supply, capital accumulation, and real output. This component is entirely absent in the MIU model.
- The Bailey triangle is therefore a **lower bound** on the true welfare cost of inflation when the CIA friction is present. Computing the full cost requires solving for the general-equilibrium response of $(c^\star, n^\star, k^\star)$ to changes in μ.
- At the **Friedman Rule** ($i^\star = 0$): both components vanish simultaneously, restoring the first-best real allocation and eliminating the liquidity distortion in one shot.

The OLG Model: A Qualitatively Different Cost

- In MIU and CIA, money is valued because it provides liquidity *within* a period. In OLG, money's role is fundamentally different: it is a **store of value** that allows the young to transfer resources to their old selves by trading with the next generation.

- The welfare cost of inflation is therefore not about reduced money holdings in a portfolio sense. It is about the deterioration of money as an *intergenerational trading mechanism*.

- Recall that in OLG the real return on money satisfies:

$$1 + r_{t+1} = \frac{p_t}{p_{t+1}} = \frac{1}{1 + \pi_{t+1}} \approx 1 - \pi_{t+1}.$$

  When inflation is positive, the real return on money is **negative**: a unit of goods saved today buys less than a unit when the saver is old.

- This erodes the incentive to save in money, and with it the very foundation of intergenerational trade. The Bailey triangle **does not apply** here: there is no money demand curve in the MIU/CIA sense, and the welfare cost operates through a completely different channel.

The OLG Model: Three Distinct Welfare Effects

- In the OLG framework, inflation generates three welfare effects that are absent in MIU and CIA:

  1. **Distributional effect**: inflation transfers wealth from the *old* - who hold money and receive less purchasing power when they sell it - to the *young*, who receive lump-sum seigniorage transfers. Inflation is a tax on the old and a subsidy to the young.
  2. **Efficiency effect**: from the Balasko-Shell theorem, the monetary equilibrium is Pareto efficient only if $\sum_t p_t < \infty$, which requires the real return on money to be sufficiently high (i.e., inflation sufficiently low). Positive inflation pushes prices upward, potentially violating this condition and rendering the equilibrium *inefficient*.
  3. **Collapse risk**: if inflation is high enough that the real return on money falls below autarky levels, the young generation prefers not to trade at all. The monetary equilibrium *unravels entirely* - society reverts to autarky, a catastrophic and non-marginal welfare cost that no triangle can measure.

- These effects are cumulative: mild inflation causes a distributional distortion; moderate inflation causes inefficiency; extreme inflation destroys money altogether.

Other Inflation Costs

- The previous measures only the *direct* welfare cost from reduced money holdings. Several important extensions push the true cost higher:
  - **Shoe-leather costs**: households make more frequent trips to the bank to convert interest-bearing assets into cash. These trips consume real time and resources. Higher inflation $\rightarrow$ more trips $\rightarrow$ larger resource waste. The Bailey triangle is a reduced-form proxy for these costs.
  - **Menu costs**: under higher inflation, firms must update prices more frequently. Price adjustment is costly (reprinting menus, updating software, negotiating contracts). These real resource costs are absent from the Bailey triangle.
  - **Tax-system interactions**: most tax systems are not fully indexed to inflation. Higher inflation erodes the real value of depreciation allowances, generates capital gains on nominal assets, and causes bracket creep in progressive income taxes - all of which distort investment and labor supply beyond the money demand channel.
  - **Distributional effects**: inflation is a regressive tax. Poorer households hold a larger share of their wealth in cash and have less access to inflation-hedging financial instruments. The Bailey triangle, which measures average welfare, misses the disproportionate burden on the poor.

What the Models Tell Us: A Comparison

- The three models agree on the **direction**: inflation is costly, and the Friedman Rule ($i^\star = 0$) is optimal. But they disagree on the **mechanism** and **magnitude**:

	MIU	**CIA**	**OLG**
Channel	Liquidity	Labor/output	Store of value
Superneutrality?	Yes	No	No
Bailey triangle?	Exact	Lower bound	Not applicable
Real output effect?	No	Yes	Yes
Distributional?	No	No	Yes

- Each row reflects a substantive difference in model structure: MIU treats money as yielding direct utility; CIA imposes a liquidity constraint that distorts the labor margin; OLG assigns money the role of an intergenerational transfer device.

What the Models Tell Us: Key Lessons

- **Model choice matters for policy.** A central bank that uses the MIU model to calibrate the welfare cost of inflation will systematically *underestimate* it. Real output distortions (CIA) and distributional effects (OLG) are invisible in MIU - not because they are small, but because the model was not built to see them.

- **The Bailey triangle is a lower bound, not the full story.** Even within MIU and CIA, Lucas (2000) yields estimates between 0.45% and 0.9% of GDP at 10% inflation. CIA adds a labor-market distortion above the triangle; OLG adds distributional transfers and Pareto inefficiencies for which no triangle exists.

- **Beware of catastrophic, non-marginal costs.** The OLG model issues a warning that MIU and CIA cannot: if inflation erodes the real return on money below autarky levels, the welfare cost is not a triangle - it is the complete unraveling of monetary exchange. Hyperinflations are not large triangles; they are equilibrium collapses.

Outline

1. A Primer on Monetary Economics
2. Foundations of Money
3. The Welfare Cost of Inflation
4. Deeper Foundations of Money
5. Quantity Theory, Expectations, and Policy
6. Time Inconsistency and Central Bank Design
7. Monetary-Fiscal Interactions
8. The Role of Private Banks

Why Do We Need a Deeper Foundation?

- **MIU** puts money in the utility function - a shortcut for "money is useful," not an explanation of *why*.
- **CIA** imposes a cash-in-advance constraint by assumption - silent on why credit cannot simply replace cash.
- **OLG** gives a real reason - non-overlapping generations - but one tied to a very particular demographic structure.
- In every case, money was *assumed*, then its consequences derived. None asks the more primitive question: when does an object like money need to exist at all?

Why Do We Need a Deeper Foundation?

- This section answers that question from two complementary angles:
  - **Kocherlakota (1998)**: in what environments is money *essential* - strictly better than the best feasible credit arrangement?
  - **Kiyotaki and Wright (1989, 1993)**: how does a particular object - possibly one with no intrinsic value - come to be used as money through decentralized, self-interested choices?
- Keep the OLG model's "money is a bubble" result in mind (Samuelson's theorem: with infinitely-lived agents and no frictions, money is never valued) - what follows explains which frictions do the work.

Money is Memory: The Core Idea

- The MIU, CIA, and OLG models *assume* that money is used - they do not explain why. Kocherlakota (1998, *Journal of Economic Theory*) asks a deeper question: in what environments is money **essential**?

- The core insight: money is a substitute for **record-keeping**. Consider an economy where a social planner can observe and store every agent's full trading history - perfect **memory**. Then credit arrangements can replicate all allocations achievable with money: a seller who just produced for a stranger can be credibly compensated later, because everyone's obligations are publicly recorded.

- **Formal equivalence (Kocherlakota 1998)**: any equilibrium allocation supportable with money can also be supported by a credit economy with perfect memory, and vice versa. **Money and memory are perfect substitutes.**

- The intuition: carrying money is a portable, self-enforcing record of past economic contributions. When you hand over currency, you are conveying information - "I produced something for someone, somewhere" - without requiring a centralized registry or trust between strangers.

Money is Memory: Lessons and Implications

- Money is **essential** if and only if two frictions are jointly present:
  - **Anonymity**: trading partners cannot observe each other's histories - credit is impossible because debts cannot be verified or enforced bilaterally.
  - **Limited commitment**: agents can defect from future obligations. Without enforcement, promises are not credible.

- Remove either friction and money becomes redundant: with perfect memory, reputation and informal credit replace money (village economies, close-knit networks); with full commitment, complex contingent contracts suffice.

- **A contemporary implication**: digital payment systems and central bank digital currencies (CBDCs) create permanent, observable transaction records - reducing anonymity. As record-keeping improves, the transaction demand for cash may fall, not because money loses value but because the informational role it plays becomes redundant.

- **Limitation**: Kocherlakota establishes an *equivalence*, not an emergence theory. The model takes the friction as given and shows money can solve it - but does not explain why bilateral exchange with anonymous strangers is the dominant market structure.

Debate: Search Frictions or Legal Restrictions?

- Kocherlakota treats anonymity and limited commitment as facts about the trading environment. **Neil Wallace** (whom we meet again via Sargent-Wallace in the Monetary-Fiscal Interactions section) argues this is too generous to markets: real economies have institutions - banks, clearinghouses, credit bureaus - that could plausibly solve record-keeping, if allowed to.

- His **legal restrictions theory** (1983): fiat money's use is sustained not purely by informational frictions but by *legal* ones - reserve requirements, restrictions on private notes, legal tender laws. Without them, private interest-bearing IOUs would dominate fiat currency and drive it out of circulation.

Debate: Search Frictions or Legal Restrictions?

- This is the **OLG** puzzle again: money pays no interest yet is held anyway - a "bubble." Search theory says trade frictions make holding it rational; legal restrictions theory says regulation makes it necessary.

- **Discussion**: the two are not mutually exclusive - one explains why dominated money *can* be valued, the other why higher-return competitors do not simply displace it. Which feels more relevant for cash today? For a CBDC?

The Double Coincidence Problem

- Barter - direct exchange of goods without money - requires a **double coincidence of wants**: agent A must want what B has *and* B must want what A has, simultaneously and in matching quantities.

- Jevons (1875): *"The first difficulty in barter is to find two persons whose disposable possessions mutually suit each other's wants. This is called the double coincidence, which is clearly less likely to occur the greater the number of commodities, and the more completely their production is specialized."* In an economy with N goods, the probability of a double coincidence between two randomly matched agents is approximately $1/N$ - negligible in large, specialized economies.

- A simple example: a baker needs shoes, a cobbler needs meat, and a butcher needs bread. No bilateral barter is possible. Feasible solutions - chains of trades, storage, IOUs - all require either extreme coordination costs or enforcement institutions.

- **Money separates the two legs of every transaction**: sell to anyone, buy from anyone. The medium of exchange breaks the requirement for simultaneity and bilateral reciprocity. The question is: how does money emerge from individual incentives? This is what Kiyotaki and Wright (1989, 1993) formalize.

Kiyotaki-Wright: The Environment

- There are three types of agents - call them Type 1, Type 2, and Type 3 - in equal proportions. Each type is specialized:

Agent	Produces	Wants to consume
Type 1	Good 1	Good 2
Type 2	Good 2	Good 3
Type 3	Good 3	Good 1

- Each agent holds **one unit of a good** at a time. Trade is voluntary and bilateral: in each period, agents are randomly matched in pairs and decide whether to swap goods. Consuming your desired good gives utility $u > 0$; after consuming, you immediately produce one unit of your own good.

- There are **no double coincidences**: no two agents simultaneously want what the other has. Check the table - Type 1 wants good 2 (Type 2's good), but Type 2 wants good 3, not good 1. **Pure barter is impossible.**

Kiyotaki-Wright: Why Barter Fails

- Consider every possible meeting and check whether both agents want to trade:

Meeting	**Who wants to trade?**	**Barter possible?**
Type 1 meets Type 2	Only Type 1 (wants good 2)	No
Type 2 meets Type 3	Only Type 2 (wants good 3)	No
Type 3 meets Type 1	Only Type 3 (wants good 1)	No

- In every possible meeting, one agent wants what the other has - but the other does not reciprocate. Trade requires *mutual* consent, so **no trade ever takes place under pure barter**. Everyone holds their own good forever. This is the double coincidence problem made precise.

- Yet clearly there are **gains from trade**: if Type 1 could somehow get good 2, she would enjoy utility u. The problem is not the absence of gains - it is the absence of a mechanism to realize them. This is exactly the problem that a **medium of exchange** solves.

Kiyotaki-Wright: The Logic of Indirect Exchange

- Suppose Type 1 (holding good 1) meets Type 3 (holding good 3). Type 3 *wants* good 1 and offers good 3 in exchange. Type 1 does *not* consume good 3 - but should she accept it anyway?
- Consider Type 1's reasoning: *"Good 3 is exactly what Type 2 wants. If I accept good 3 now, the next time I meet a Type 2 agent, I can trade good 3 for good 2 - which is what I actually want."*
- This is **indirect exchange**: accept a good you do not consume because you expect to pass it on to someone who does. Good 3 is acting as a **medium of exchange** - a stepping stone between production and consumption.
- Notice the crucial feature: Type 1's willingness to accept good 3 depends entirely on her expectation that **Type 2 will also accept good 3**. If Type 2 refuses good 3, then holding good 3 is useless and Type 1 should not accept it. **The value of a medium of exchange is inherently about shared expectations.**

Kiyotaki-Wright: Value Functions

- Let r be the discount rate and α the probability of meeting each type per unit of time. Consider a Type 1 agent who follows the strategy of *accepting good 3 as money*. At any point in time she is in one of two states:
  - **Holding good 1** (her own good): she is waiting to meet a Type 3 agent who will offer good 3 in exchange. Value: V_1.
  - **Holding good 3** (the stepping stone): she is waiting to meet a Type 2 agent who will accept good 3 and give her good 2. Value: V_3.
- The flow **Bellman equations** are:

$$r\,V_1 = \alpha(V_3 - V_1)$$
$$r\,V_3 = \alpha(u + V_1 - V_3)$$

- Reading these equations: the left side is the flow cost of waiting (discounting). The right side is the expected capital gain from the next meeting. For V_1: with probability α she meets Type 3 and jumps from state V_1 to state V_3. For V_3: with probability α she meets Type 2, consumes (gains u), and resets to V_1.

Kiyotaki-Wright: The Commodity Money Equilibrium

- Is it **individually rational** for Type 1 to accept good 3? Yes - as long as $V_3 > V_1$. Holding the liquid good (good 3) gets the agent closer to her desired consumption than holding her own good, which no one wants to buy.

- Adding the two Bellman equations gives $r(V_1 + V_3) = \alpha(u + 2V_1 - 2V_1)$... subtracting the first from the second:

$$r(V_3 - V_1) = \alpha(u + V_1 - V_3) - \alpha(V_3 - V_1) = \alpha\, u - 2\alpha(V_3 - V_1)$$

$$\implies \quad V_3 - V_1 = \frac{\alpha\, u}{r + 2\alpha} > 0$$

  $V_3 > V_1$ **always holds**: holding the medium of exchange is strictly more valuable than holding your own good. Accepting good 3 is always individually rational.

- This is a **Nash equilibrium**: each agent's strategy (accept the stepping-stone good) is a best response to the others' strategies. Good 3 functions as **commodity money** - it has direct consumption value for Type 2 agents, which is what makes it liquid and trustworthy for everyone else.

Kiyotaki-Wright: From Commodity Money to Fiat Money

- In the commodity money equilibrium, good 3 works as a medium of exchange because it has *intrinsic* value: Type 2 agents genuinely want to consume it. That direct value is what makes it credible as a stepping stone.

- Now ask: does the medium of exchange *need* to have intrinsic value? Consider replacing good 3 with an object that **no one wants to consume** - a token, a coin, a piece of paper. Call it **fiat money** m. Its direct utility is zero: $u_m = 0$.

- The Bellman equations for a Type 1 agent who accepts fiat money are structurally identical:

$$r\,V_1 = \alpha(V_m - V_1) \qquad r\,V_m = \alpha(u + V_1 - V_m)$$

  The second equation looks the same as before - but now the u on the right is the utility from eventually consuming *good 2*, not from consuming m itself. Fiat money is purely a vehicle; all the value comes from what it buys.

- Whether $V_m > V_1$ holds - whether accepting fiat money is rational - depends entirely on whether other agents also accept it. This is the key difference from commodity money.

Kiyotaki-Wright: Multiple Equilibria and Social Conventions

- **Accepting fiat money is individually rational if and only if others accept it too.** This self-referential logic generates **two equilibria**:
  - **Monetary equilibrium**: all agents accept fiat money $\Rightarrow V_m > V_1 \Rightarrow$ accepting is a best response $\Rightarrow$ belief is self-fulfilling. Fiat money circulates; everyone trades; the economy escapes autarky.
  - **Non-monetary equilibrium**: no agent accepts fiat money $\Rightarrow V_m = V_1$ (holding m is useless) $\Rightarrow$ rejecting is a best response $\Rightarrow$ belief is also self-fulfilling. Fiat money is worthless; the economy is stuck in autarky.

- Fiat money has value **purely because of coordinated belief**. It is a **social convention**: not a technological necessity, not a government decree, but a shared equilibrium expectation. History is full of examples - shells, cigarettes in POW camps, Bitcoin.

- **The fragility of fiat money**: unlike commodity money, the monetary equilibrium under fiat money can unravel. If enough agents lose confidence and stop accepting m, the expectation of acceptance collapses - and with it, the value of money itself. This is the microeconomic foundation of a monetary crisis.

Case Study: Money Without a State

- **POW camps (Radford, 1945)**: WWII Red Cross parcels held standardized goods but no currency. Cigarettes spontaneously became the medium of exchange - prices quoted in cigarettes, with "inflation" when new shipments arrived and contractions after raids destroyed stockpiles. Kiyotaki-Wright in miniature: even non-smokers held cigarettes purely for their acceptability.

- **Somalia, 1991-2011**: after the central government collapsed, no new shillings were printed - yet the old, unbacked banknotes kept circulating and holding value for two decades, sustained purely by coordinated belief.

Case Study: Money Without a State

- **The Iraqi Swiss dinar (1990s)**: after the Gulf War, Iraq introduced a new dinar, but the Kurdish north kept using the old, no-longer-printed "Swiss dinar" - which traded at a *premium* to the new currency, precisely because its supply could not be inflated further.

- **The common lesson**: fiat value does not require a functioning state - it requires *coordinated belief in future acceptability*. States are usually good at engineering that belief (legal tender laws, taxation), but as these examples show, they are neither necessary nor sufficient.

Kiyotaki-Wright: Key Lessons

- **Money is endogenous.** Unlike MIU and CIA, Kiyotaki-Wright does not assume money. Money emerges from the strategic behavior of self-interested agents facing a search friction. Any object - a commodity, a coin, a piece of paper - can serve as money if agents coordinate on accepting it.

- **Liquidity is about acceptability, not intrinsic value.** A good is liquid if it is widely accepted in trade. Commodity money (good 3) is liquid because it has intrinsic value to some agents, which gives others a reason to hold it. Fiat money is liquid for a more fragile reason: everyone accepts it because everyone accepts it.

- **Money solves a coordination problem.** The non-monetary equilibrium (autarky) is also an equilibrium - and it is worse for everyone. This means that the existence of a monetary system is not guaranteed: it requires *coordination on the right equilibrium*. Institutions, history, and trust play an essential role.

- **A lesson for policy.** Hyperinflations are not just large triangles (as the Bailey framework suggests). They can be *equilibrium switches*: if inflation erodes trust in fiat money sufficiently, agents may stop accepting it altogether - the monetary equilibrium collapses into autarky. The OLG model pointed to this risk; Kiyotaki-Wright makes the mechanism precise.

Why Does the Economy Settle on One Money?

- Kiyotaki-Wright *assumed* a single candidate medium of exchange (good 3). But nothing rules out multiple monetary equilibria, with different goods playing the role of money. Why do economies tend to converge on *one* widely accepted money?

- **Network externalities**: accepting an object as money is more valuable the more *other* agents also accept it - more potential trading partners at every future meeting. The same logic makes a single social media platform, or a single language, dominant even absent intrinsic superiority.

Why Does the Economy Settle on One Money?

- This is also why fiat currencies have **strong incumbency advantages**: a new currency (foreign, or a private digital token) must overcome a coordination problem - getting agents to jointly switch expectations - not just win individual willingness to accept it.

- **Discussion**: this also explains *dollarization* - once agents in a high-inflation economy coordinate on a foreign currency, network externalities can lock it in even after domestic inflation is tamed. Multiplicity of equilibria has first-order policy consequences.

Beyond Kiyotaki-Wright: New Monetarist Economics

- Basic Kiyotaki-Wright has two uncomfortable features for quantitative work: goods are indivisible and agents hold at most one unit of money - no sense of "how much" money an agent holds, and no role for prices to adjust continuously.

- **Lagos and Wright (2005)** fix this by alternating, each period, a decentralized search market (like Kiyotaki-Wright) with a **centralized Walrasian market** where agents freely trade and rebalance money holdings - restoring tractability without losing the search-theoretic foundation for why money is used.

Beyond Kiyotaki-Wright: New Monetarist Economics

- This framework - **New Monetarist Economics** - is the workhorse of modern search-theoretic monetary theory: it reintroduces continuous money, studies the welfare cost of inflation with explicit microfoundations, and analyzes policy under genuine trading frictions.

- The takeaway, without the algebra: everything from Kiyotaki-Wright about *why* money is valued survives; New Monetarist models simply make it quantitatively usable, connecting search-theoretic money back to the same policy questions MIU and CIA were built to answer.

Applying the Lenses: Is Bitcoin Money?

- Bitcoin and similar cryptocurrencies are a natural test case for the two frameworks in this section - a discussion exercise, not a settled answer.

- **Through the Kocherlakota lens**: the blockchain is a perfect, public, permanent record of every transaction - the opposite of anonymity. If histories are perfectly observable, credit arrangements should replicate anything Bitcoin achieves. Its "essentiality," if any, comes from distrust in centralized record-keepers, not from anonymity.

Applying the Lenses: Is Bitcoin Money?

- **Through the Kiyotaki-Wright lens**: to function as a medium of exchange, merchants must expect others to accept it later - the same coordinated-belief logic. In practice, adoption for actual transactions is limited; Bitcoin is held mostly as a speculative store of value, rarely passed along as a stepping-stone.

- **Discussion**: does Bitcoin fail to be "money" from insufficient coordination (a non-monetary equilibrium), or because it solves a different problem altogether (censorship-resistant store of value) than the double-coincidence problem this section addresses?

Outline

1. A Primer on Monetary Economics
2. Foundations of Money
3. The Welfare Cost of Inflation
4. Deeper Foundations of Money
5. Quantity Theory, Expectations, and Policy
6. Time Inconsistency and Central Bank Design
7. Monetary-Fiscal Interactions
8. The Role of Private Banks

The Quantity Theory of Money: The Equation of Exchange

- The oldest and most influential proposition in monetary economics is the **Quantity Theory of Money**, formalized by Irving Fisher (1911) and later by the Cambridge economists. Its foundation is the **equation of exchange**:

$$M_t V_t = P_t Y_t$$

where M_t is the nominal money supply, V_t is the **velocity of money** (the average number of times each dollar is spent on final goods per period), P_t is the price level, and Y_t is real output. This is an accounting identity - it holds by definition.

- The theory arises from two classical assumptions:
  - **Velocity is stable**: $V_t \approx \bar{V}$. The payment habits of households and the structure of financial markets change slowly.
  - **Real output is determined by the real side of the economy**: technology, labor, and capital - not by money. This is the classical dichotomy.

- Under these two assumptions, the equation of exchange becomes a theory: $M_t \bar{V} = P_t Y_t^{\star}$, so the price level is proportional to the money supply. **More money means higher prices, one-for-one.**

The Quantity Theory of Money: The Key Prediction

- Taking logarithms of $MV = PY$ gives the **growth-rate version**:

$$\mu_t + \hat{v}_t = \pi_t + g_t$$

where μ_t is money growth, $\hat{v}_t$ is velocity growth, π_t is inflation, and g_t is real output growth.

- If velocity is stable ($\hat{v}_t = 0$) and real output grows at its long-run trend $\bar{g}$ (determined independently of money), then:

$$\boxed{\pi_t = \mu_t - \bar{g}}$$

Inflation equals money growth minus real growth. A central bank that grows the money supply faster than the economy grows real output will generate inflation - dollar for dollar above trend.

- This is the substance of Friedman's (1963) famous dictum: *"Inflation is always and everywhere a monetary phenomenon."* Not a claim that every price change is caused by money, but that sustained, large inflation is always traceable to sustained, excessive money creation.

- **Note**: the theory is about the *long run*. In the short run, velocity fluctuates and real output responds to shocks, so the relationship between M and P is noisy. But over long horizons and large money growth episodes, it holds with remarkable precision.

The Quantity Theory of Money: Empirical Evidence

- **Cross-country evidence**: McCandless and Weber (1995) examined 110 countries over a 30-year period and found a near one-for-one relationship between average money growth and average inflation - exactly as the quantity theory predicts. Countries with moderate monetary expansion had moderate inflation; countries that printed money aggressively had high inflation, without exception.
- **Hyperinflations**: the most compelling evidence comes from extreme episodes:
  - **Germany (1921–1923)**: the Reichsbank financed government deficits by printing money. By November 1923, prices were rising at over 30,000% per *month*. Workers were paid twice a day so wages could be spent before they lost value.
  - **Zimbabwe (2007–2008)**: peak monthly inflation exceeded 79 billion percent. The central bank printed 100-trillion-dollar notes. The currency was eventually abandoned.
  - **Venezuela (2010s–2020s)**: annual inflation exceeded 1,000,000% at its peak, driven by monetization of fiscal deficits.
- **The short-run caveat**: over business cycle frequencies, the money-inflation link is weaker. Velocity is not perfectly stable, and short-run real effects of money exist - as Lucas (1972) showed.

Lucas (1972): Money, Information, and the Islands Economy

- The quantity theory tells us that money is neutral in the long run. But we observe that monetary shocks seem to have real effects in the short run - recessions follow tight monetary policy, and booms follow expansions. How can we reconcile these facts?

- Lucas (1972, *Journal of Economic Theory*) provides the answer through the concept of **imperfect information**. The economy consists of many separate **islands**, each with its own local market. Agents live on their island and observe only the **local price** of their good - they do not directly observe the aggregate price level P_t or the money supply M_t.

- When a producer on island i sees the local price p_{it} rise, she faces a **signal extraction problem**:
  - **Interpretation A**: the demand for *my* good has risen - a real, relative price increase. Rational response: produce more.
  - **Interpretation B**: the *aggregate* price level has risen due to money creation - an inflationary nominal increase. Rational response: do nothing (relative prices have not changed).

- The producer cannot perfectly distinguish between these two cases from local price information alone. She uses statistical inference to form the best possible estimate - but she will *always* partially attribute aggregate inflation to a local real shock.

Lucas (1972): Real Effects of Monetary Surprises

- The signal extraction problem generates a crucial asymmetry between **anticipated** and **unanticipated** monetary policy:
  - **Anticipated money growth**: if agents know the central bank will increase M_t by 10%, they expect prices everywhere to rise by 10%. When they see local prices rise, they correctly attribute it to aggregate inflation and do not change production. **Money is neutral**.
  - **Unanticipated money growth**: a surprise increase in M_t raises all prices, but agents don't know this. They partially interpret rising local prices as a local demand increase and *do* expand production. Real output rises - temporarily.
- This gives rise to the **Lucas aggregate supply curve**:

$$y_t = y^\star + b(\pi_t - \pi_t^e), \qquad b > 0$$

  Output deviates from its natural level $y^\star$ only when inflation surprises agents - when $\pi_t \neq \pi_t^e$. The parameter b reflects how much producers respond to price signals.
- **The key lesson**: systematic monetary policy cannot be used to permanently raise output above $y^\star$. Only *surprises* generate real effects, and rational agents will not be fooled systematically.

The Phillips Curve: A Stable Trade-Off?

- In 1958, Alban William Phillips published a study of British data from 1861 to 1957 and documented a remarkably stable negative relationship between **wage inflation** and the **unemployment rate**. Samuelson and Solow (1960) extended this to US data and reframed it in terms of price inflation, producing the **Phillips curve** as we know it: when unemployment is low, inflation tends to be high, and vice versa.

- The empirical relationship was striking. For the United States in the 1950s and 1960s, the data fit a smooth downward-sloping curve with remarkable consistency.

- The Phillips curve became the centerpiece of Keynesian macroeconomic policy in the 1960s. It seemed to say: *if you want to reduce unemployment, just accept a bit more inflation*. The trade-off appeared permanent, exploitable, and quantifiable.

The Expectations-Augmented Phillips Curve

- Friedman (1968) and Phelps (1968) independently warned that the apparent trade-off was an illusion: workers and firms care about **real wages**, not nominal wages. Any attempt to reduce unemployment by raising inflation will be met with revised wage demands once workers update their inflation expectations.

- The correct relationship is the **expectations-augmented Phillips curve**:

$$\pi_t = \pi_t^e - \alpha(u_t - u^\star), \qquad \alpha > 0$$

  where $u^\star$ is the **natural rate of unemployment** - determined by real factors (search frictions, skill mismatch, labor market institutions), not by monetary policy.

- This equation has two very different implications depending on the time horizon:
  - **Short run** (π_t^e fixed): a trade-off exists. Surprise inflation lowers unemployment below $u^\star$.
  - **Long run** ($\pi_t^e = \pi_t$): there is no trade-off. Setting $\pi_t = \pi_t^e$ gives $u_t = u^\star$ regardless of the inflation rate. The **long-run Phillips curve is vertical** at $u^\star$.

- The implication: monetary policy cannot permanently reduce unemployment below $u^\star$. Any attempt to do so results in **ever-accelerating inflation** as expectations continuously chase inflation.

The Breakdown: Stagflation and the End of the Trade-Off

- The 1970s provided a definitive empirical verdict. Following expansionary monetary and fiscal policies aimed at keeping unemployment low, the United States - and much of the developed world - experienced **stagflation**: simultaneously high inflation *and* high unemployment. This was impossible under the original Phillips curve and precisely what Friedman and Phelps had predicted.
- What happened? Policymakers in the 1960s ran expansionary policy to exploit the apparent trade-off. Workers and firms gradually learned to anticipate higher inflation and built it into wage contracts. Each round of inflation raised expectations, which fed into the next round - an **inflationary spiral**. By the late 1970s, US annual inflation exceeded 13% while unemployment remained elevated.
- The Volcker disinflation (1979–1983) restored the lesson in reverse: breaking inflationary expectations required a severe recession (unemployment peaked at nearly 11%), confirming that the short-run trade-off is real but costly to exploit - and ultimately self-defeating.
- **What the episode teaches**: the original Phillips curve was a *statistical regularity*, not a structural relationship. It held as long as inflation expectations were anchored at low levels. The moment policy tried to systematically exploit it, agents updated their expectations and the curve shifted.

The Rational Expectations Revolution

- Before the 1970s, most macroeconomic models used **adaptive expectations**: agents form expectations by looking backward, extrapolating from past inflation. For example: $\pi_t^e = \pi_{t-1}$, or a weighted average of past inflation rates. This is simple but has a fatal flaw - it implies that agents make *systematic, predictable errors* that they never correct.

- Muth (1961) proposed **rational expectations**: agents use all available information efficiently, including knowledge of the model itself. Formally:

$$\pi_t^e = \mathbb{E}[\pi_t \mid \Omega_{t-1}]$$

  where Ω_{t-1} is the agent's information set at time $t-1$ - everything known before π_t is realized. Rational expectations does not mean perfect foresight: agents can still be wrong, but their errors are *unpredictable* (mean zero, uncorrelated with anything in Ω_{t-1}).

- Lucas, Sargent, and Wallace applied rational expectations to macroeconomics in the 1970s, producing a revolution in how economists model policy. The core insight: **if agents understand the policy rule, they will incorporate it into their expectations before the policy takes effect** - neutralizing any intended real impact. You cannot systematically fool rational agents.

Policy Implications of Rational Expectations

- Combining rational expectations with the Lucas aggregate supply curve $y_t = y^\star + b(\pi_t - \pi_t^e)$ yields the **policy ineffectiveness proposition** (Sargent and Wallace, 1975):
  - Suppose the central bank follows a systematic rule: "whenever unemployment rises, increase money growth by x%." Rational agents know this rule, anticipate the money growth, and immediately revise π_t^e upward. By the time the policy acts, $\pi_t - \pi_t^e = 0$ - **no real effect**.
  - Only a **surprise** departure from the known rule ($\pi_t \neq \pi_t^e$) generates a real output response. But the central bank cannot plan a surprise - by definition, a planned surprise is not a surprise.
- **What monetary policy can do**:
  - Anchor inflation expectations by committing to a credible nominal target.
  - Reduce the variance of surprise inflation, which itself has real costs (uncertainty, misallocation).
  - Respond to shocks that *agents also cannot observe*, preserving some scope for stabilization policy.
- The rational expectations revolution did not render monetary policy irrelevant - but it permanently changed the questions we ask: not "how much output can we buy with inflation?" but "how do we build credible institutions that anchor expectations?"

The Lucas Critique (1976)

- In a landmark paper - "Econometric Policy Evaluation: A Critique" - Lucas (1976) identified a fundamental flaw in the way macroeconomists used empirical models for policy analysis.

- The **standard approach** of the era: estimate a large econometric model on historical data, capturing the statistical relationships between variables (output, inflation, unemployment, interest rates). Then simulate the model to predict what would happen under a proposed new policy.

- **The problem**: the estimated statistical relationships are not structural. They are the *reduced-form outcomes* of agents' optimal decisions under the current policy regime. When the policy regime changes, agents *change their behavior* - and the estimated parameters shift with them.

- Lucas's argument: *"Given that the structure of an econometric model consists of optimal decision rules of economic agents, and that optimal decision rules vary systematically with changes in the structure of series relevant to the decision maker, it follows that any change in policy will systematically alter the structure of econometric models."*

- In short: **you cannot use a model estimated under one policy regime to predict behavior under a different policy regime** - because the model itself will change when the policy changes.

The Lucas Critique: The Phillips Curve and Its Legacy

- The **Phillips curve** is the canonical example. When governments tried to exploit the curve, workers and firms anticipated the higher inflation and demanded higher wages preemptively. The Phillips curve *shifted*: both unemployment and inflation rose together - the stagflation of the 1970s. The estimated relationship was not structural; it was an artifact of the low-inflation policy regime.

- **The correct lesson**: to evaluate any policy change, we need models built on **deep parameters** - preferences, technology, and endowments - that are invariant to the policy being studied. These parameters describe *how people fundamentally behave*, not how they behave under a particular regime.

- **The legacy of the Lucas Critique**: it is the intellectual foundation of modern quantitative macroeconomics. Dynamic Stochastic General Equilibrium (DSGE) models - the workhorse of central banks worldwide - are built precisely to have structural, regime-invariant parameters. Every model in this course is a step in that direction: optimization-based, with explicit preferences and technology, immune to the critique by construction.

Outline

1. A Primer on Monetary Economics
2. Foundations of Money
3. The Welfare Cost of Inflation
4. Deeper Foundations of Money
5. Quantity Theory, Expectations, and Policy
6. Time Inconsistency and Central Bank Design
7. Monetary-Fiscal Interactions
8. The Role of Private Banks

Kydland-Prescott (1977): The Time Inconsistency Problem

- Kydland and Prescott (1977, *Journal of Political Economy*) identified a fundamental difficulty in the conduct of optimal economic policy: a plan that is *optimal today* may no longer be optimal to carry out *tomorrow*, once private agents have already made decisions based on the announced plan. When this happens, the policymaker has an incentive to deviate - and rational agents will anticipate this.

- A simple non-monetary example makes the point. Suppose a government announces it will never bail out failing banks (*no bailout* policy). Banks, believing the announcement, exercise caution. But once a large bank is on the verge of collapse, the government faces a new calculation: the cost of following through (systemic crisis) may exceed the cost of deviating (one bailout). The policy is **time-inconsistent**: optimal ex-ante (before banks decide), but not optimal ex-post (after the crisis materializes).

- In monetary policy, the same logic applies. The central bank would like to announce low inflation to anchor expectations. But once workers and firms have set wages based on that announcement - *once expectations are locked in* - the central bank faces a temptation: generate surprise inflation to push output above its natural level. Rational agents foresee this temptation and do not fully believe the announcement.

Kydland-Prescott: Why the Temptation Exists

- In monetary policy, the time inconsistency temptation arises from the **natural level of output being inefficiently low**. Due to taxes, monopoly markups, or labor market frictions, actual output $y^\star$ falls short of the socially optimal level. The central bank would like to close this gap.

- The interaction between the central bank and the private sector has a sequential structure:

  1. The central bank **announces** an inflation target π^a (say, zero).
  2. Workers and firms **set wages and prices** based on the announcement: $\pi^e = \pi^a$. Wage contracts are signed; expectations are locked in.
  3. The central bank **chooses actual inflation** π, taking π^e as given. Given the expectation-augmented Phillips curve $y = y^\star + \alpha(\pi - \pi^e)$, the central bank can now raise output by setting $\pi > \pi^e$. The benefit (higher output) is immediate; the cost (higher inflation) is shared.

- If the central bank optimizes at step 3, it will *always* deviate from its announcement. But rational agents at step 2 know this - they do not believe the zero-inflation announcement, set $\pi^e > 0$, and the economy ends up with high inflation *and* no output gain. This is the **inflation bias**: the equilibrium under discretion is worse than what the central bank itself would prefer if it could commit.

Barro-Gordon (1983): The Model

- Barro and Gordon (1983) formalize the time inconsistency problem in a tractable model. The economy has two building blocks:

- **Aggregate supply** (expectations-augmented Phillips curve):

$$y_t = y^\star + \alpha(\pi_t - \pi_t^e), \qquad \alpha > 0.$$

- **Central bank loss function**: the central bank minimizes a quadratic loss in inflation and deviations of output from a target level $y^\star + k$, where $k > 0$ captures the output distortion:

$$L = \frac{1}{2}\left[\pi^2 + \lambda(y - y^\star - k)^2\right], \qquad \lambda > 0$$

  The term λ reflects the relative weight placed on output stabilization versus price stability. The distortion $k > 0$ is the source of all the trouble: the central bank wants output *above* the natural level, which is exactly what creates the temptation to inflate.

- Without the distortion ($k = 0$), the economy would achieve $\pi = 0$, $y = y^\star$ - the first best - even under discretion. **Time inconsistency is a problem only because the policymaker's goal is distorted.**

Barro-Gordon: The Inflation Bias Under Discretion

- Under **discretion**, the central bank takes private expectations π^e as given and minimizes L over π. Substituting the supply curve $y = y^\star + \alpha(\pi - \pi^e)$:

$$L = \frac{1}{2}\left[\pi^2 + \lambda(\alpha(\pi - \pi^e) - k)^2\right]$$

- First-order condition:

$$\pi + \lambda\alpha(\alpha(\pi - \pi^e) - k) = 0 \quad \Longrightarrow \quad \pi^D = \frac{\lambda\alpha^2\pi^e + \lambda\alpha k}{1 + \lambda\alpha^2}$$

- Now impose rational expectations: agents know the central bank's decision rule, so in equilibrium $\pi^e = \pi^D$. Solving for π^e:

$$\pi^e(1 + \lambda\alpha^2) = \lambda\alpha^2\pi^e + \lambda\alpha k \quad \Longrightarrow \quad \pi^e = \lambda\alpha k$$

  And therefore $\pi^D = \lambda\alpha k$. In equilibrium $\pi^e = \pi^D$, so $y = y^\star$ - **no output gain is achieved**.

- This is the **inflation bias**: $\pi^D = \lambda\alpha k > 0$. The economy suffers strictly positive inflation even though it delivers zero output benefit.

Barro-Gordon: Rules versus Discretion

- Under **commitment** (a binding rule), the central bank can credibly promise $\pi = 0$. Rational agents set $\pi^e = 0$, and $y = y^\star$. The social losses under the two regimes are:

	Inflation	**Social loss** L
Commitment (rule)	0	$\frac{1}{2}\lambda k^2$
Discretion	$\lambda\alpha k$	$\frac{1}{2}\lambda k^2\left(1+\lambda\alpha^2\right)$

- Discretion always produces a **strictly higher loss** than commitment. The inflation bias is *pure waste*: everyone would prefer the zero-inflation outcome, but no one can get there without a credible commitment device.

- **But why can't the central bank simply choose $\pi = 0$ under discretion?** If $\pi^e = 0$, the central bank would deviate to $\pi^{dev} = \frac{\lambda\alpha k}{1+\lambda\alpha^2} > 0$ to gain a short-run output boost. Knowing this, agents never set $\pi^e = 0$ in the first place. The zero-inflation outcome requires a *constraint* on the central bank's behavior - which brings us to the question of how such a constraint can be enforced.

Barro-Gordon: Reputation and Trigger Strategies

- If the central bank interacts with the public **repeatedly over time**, reputation can substitute for formal commitment. Barro and Gordon (1983) show that low inflation can be sustained as an equilibrium of an infinitely-repeated game, provided the central bank is *patient enough*.

- The **trigger strategy** works as follows. The public starts by expecting $\pi^e = 0$ (trusting the central bank). If the central bank ever deviates (inflates above 0), the public **permanently switches** to expecting $\pi^e = \lambda\alpha k$ (the discretionary level) forever after. The central bank faces a choice in every period:
  - **Cooperate** ($\pi = 0$): pay cost $L^C = \frac{1}{2}\lambda k^2$ today and every future period.
  - **Deviate**: gain a one-time benefit today (lower loss $L^{dev} < L^C$), but trigger the punishment: suffer loss $L^D > L^C$ in every future period.

- Cooperation is sustainable if the long-run cost of the punishment outweighs the short-run gain from deviating. With discount factor $\beta \in (0, 1)$, the condition is:

$$\underbrace{L^C - L^{dev}}_{\text{gain from deviating today}} \leq \frac{\beta}{1-\beta} \underbrace{\left(L^D - L^C\right)}_{\text{future cost of punishment}}$$

Rogoff (1985): The Conservative Central Banker

- One natural remedy to this problem is to change *who* conducts monetary policy. Rogoff (1985) proposes a solution based on **delegation**: appoint a central banker whose preferences differ from society's in a specific way.

- Suppose society's loss function is $L^S = \frac{1}{2}[\pi^2 + \lambda(y - y^\star - k)^2]$. Instead of a central bank with weight λ on output, appoint a **conservative central banker** with weight $\lambda^{CB} < \lambda$:

$$L^{CB} = \frac{1}{2}\left[\pi^2 + \lambda^{CB}(y - y^\star - k)^2\right], \qquad \lambda^{CB} < \lambda$$

  The conservative CB cares *more* about inflation and *less* about output than society does.

- Under discretion, the CB with weight λ^{CB} sets:

$$\pi^{CB} = \lambda^{CB}\alpha k < \lambda\alpha k = \pi^D$$

  The inflation bias is reduced. Delegating to a more inflation-averse central banker credibly signals lower inflation, which lowers expectations, which lowers actual inflation.

- This provides a theoretical foundation for the institutional design we observe in practice: central bankers are typically selected for their strong anti-inflationary reputation.

Rogoff (1985): The Stabilization Cost of Conservatism

- Delegating to an infinitely conservative central banker ($\lambda^{CB} \to 0$) would eliminate the inflation bias entirely: $\pi^{CB} \to 0$. Why not go all the way? Because extreme conservatism has a hidden cost: it destroys the central bank's ability to respond to real shocks.

- Suppose the economy is hit by a **supply shock** ε_t (e.g., an oil price spike). The supply curve becomes:

$$y_t = y^\star + \alpha(\pi_t - \pi_t^e) + \varepsilon_t$$

  A negative $\varepsilon_t < 0$ pushes output below trend. The socially optimal response is to allow *some* inflation - loosening policy cushions the output fall at an acceptable inflation cost.

- A conservative central banker, however, places very low weight on output deviations. She is reluctant to tolerate even the modest inflation that optimal stabilization requires. The result is **under-stabilization**: inflation stays low, but output falls further than society would prefer.

- This is the first horn of Rogoff's dilemma. Conservatism buys credibility in normal times, but it becomes a straitjacket when the economy needs flexibility. The trade-off is not between inflation and output on average - it is between the *level* of inflation bias and the *variance* of output around trend.

Rogoff (1985): The Optimal Degree of Conservatism

- The two effects of conservatism pull in opposite directions:

	Inflation bias	**Stabilization of shocks**
High λ^{CB} (dovish CB)	Large	Good
Low λ^{CB} (hawkish CB)	Small	Poor

- The **optimal degree of conservatism** is interior: $\lambda^{CB} \in (0, \lambda)$. Society is strictly better off appointing a central banker who is more inflation-averse than the median voter - but not one who ignores output entirely. The gain from lower inflation bias outweighs the loss in stabilization flexibility, up to a point.

- This is the central result of Rogoff (1985): **imperfect commitment through delegation dominates pure discretion, but comes at a stabilization cost**. The model does not deliver the first-best - it delivers a second-best improvement. Eliminating the last unit of inflation bias is simply too expensive in terms of lost output stabilization.

Inflation Targeting as a Policy Framework

- Framework in which the central bank commits to an inflation target over a defined time horizon, and is held accountable. New Zealand adopted it first in 1990; today it is the dominant framework.
- The framework has four defining features:
  - **Numerical target**: a publicly announced inflation objective which anchors long-run expectations.
  - **Transparency**: the central bank publishes its forecasts, models, and reasoning. Private agents can verify whether the CB is on track.
  - **Accountability**: when inflation misses the target, the CB must explain why and what it will do. Some economies require a formal letter to the government.
  - **Flexibility**: most frameworks target inflation *over the medium term*, allowing short-run deviations to stabilize output after shocks. This is **flexible inflation targeting**.
- Through the lens of our models: transparency and accountability create the reputational mechanism of Barro-Gordon; flexibility preserves the stabilization role that Rogoff's conservative CB would sacrifice.

Instrument Independence vs. Goal Independence

- A central question in central bank design is: *how much should the central bank be independent from the government?* Fischer (1995) and Debelle and Fischer (1994) draw a crucial distinction between two types of independence:
  - **Goal independence**: the central bank chooses its own objectives - what to target and how to weigh competing goals (inflation vs. output). This is a high degree of independence and is generally considered **democratically problematic**: setting macroeconomic objectives is a political decision that should be made by elected governments, not unelected technocrats.
  - **Instrument independence**: the central bank chooses *how* to achieve externally assigned objectives, without political interference in day-to-day decisions. This is widely accepted as necessary for credibility - political interference in interest rate decisions is a direct source of the inflation bias.
- Most modern central banks have **instrument independence but not goal independence**. The government sets the mandate; the central bank decides the path. In the United States, Congress established the dual mandate (price stability and maximum employment); the Federal Reserve chooses the federal funds rate. In the Eurozone, price stability is the primary mandate set by treaty; the ECB chooses instruments freely.

The Taylor Rule: A Rule that Commits the Central Bank

- Taylor (1993) proposed a simple **instrument rule** that describes how the nominal interest rate should respond to macroeconomic conditions:

$$i_t = r^\star + \pi^\star + \underbrace{1.5\,(\pi_t - \pi^\star)}_{\text{inflation gap}} + \underbrace{0.5\,(y_t - y^\star)}_{\text{output gap}}$$

  where $r^\star$ is the long-run real interest rate, $\pi^\star$ is the inflation target, and $y_t - y^\star$ is the output gap.

- The rule has two key properties:
  - **The Taylor principle**: the coefficient on inflation (1.5) exceeds 1, so the *real* interest rate $i_t - \pi_t$ rises when inflation rises. This ensures that the central bank does not accommodate inflationary shocks.
  - **Output stabilization**: the positive coefficient on the output gap (0.5) means the central bank eases when the economy is below potential and tightens when above - consistent with flexible inflation targeting.

- Taylor showed that the rule fit actual US Federal Reserve behavior from 1987 to 1992 remarkably well - even though the Fed had not explicitly adopted it.

The Taylor Rule as a Commitment Device

- Through the lens of Barro-Gordon, the Taylor rule is a practical approximation to the **optimal rule** that eliminates the inflation bias. By publicly committing to a mechanical response function, the central bank removes the discretion that generates the credibility problem: there is no occasion to "decide" whether to inflate, because the rule decides for it.

- **Limitations of the rule**:
  - **The zero lower bound (ZLB)**: when $i_t = 0$, the rule cannot prescribe negative nominal rates (with standard instruments). The rule breaks down precisely when the economy needs the most stimulus.
  - **Model uncertainty**: the output gap $y_t - y^\star$ is not observable in real time. Mismeasurement of potential output has repeatedly caused central banks to tighten or ease at the wrong moment.
  - **Financial stability**: the Taylor rule makes no reference to asset prices, credit growth, or systemic risk - factors that played a central role in the 2008 financial crisis. These limitations motivated the development of *unconventional* monetary policy tools, to which we turn next.

Unconventional Monetary Policy: The Zero Lower Bound

- Conventional monetary policy operates through the **short-term nominal interest rate**. When the central bank cuts i_t toward zero, it runs out of room: the **zero lower bound (ZLB)** is binding, and further cuts are not possible with standard instruments. The Taylor rule calls for negative rates, but cash provides a zero nominal return, placing a floor on how low rates can go.

- The ZLB became a central practical problem after the 2008 financial crisis. The Fed, the ECB, the Bank of England, and the Bank of Japan all cut rates to (or near) zero - and stayed there for years. With conventional tools exhausted, central banks turned to **unconventional monetary policy**:
  - **Forward guidance**: explicit promises about the future path of interest rates. By committing to keep rates low "for longer," the central bank lowers long-term rates even when the short rate is at zero. This works through expectations: $i_{long} \approx \mathbb{E}[i_{short}]$ over the horizon.
  - **Negative interest rates**: several central banks (ECB, SNB, Bank of Japan) pushed rates slightly below zero by charging banks for reserves. Effective, but limited in scope by the cash floor.
  - **Quantitative easing (QE)**: large-scale asset purchases designed to reduce long-term yields and stimulate the economy through channels beyond the short rate.

Quantitative Easing: Mechanisms and Assessment

- **Quantitative easing (QE)** involves the central bank purchasing large quantities of long-term assets - typically government bonds and, in some cases, mortgage-backed securities or corporate bonds - financed by creating bank reserves. The Fed's balance sheet grew from $0.9 trillion (2008) to $9 trillion (2022).
- QE operates through several channels:
  - **Portfolio balance channel**: by removing long-term bonds from the market, QE forces investors into riskier assets (equities, corporate bonds), compressing risk premia and stimulating investment and consumption.
  - **Signaling channel**: QE signals the central bank's commitment to keeping rates low for an extended period, reinforcing forward guidance and reducing long-term rates.
  - **Bank lending channel**: increased reserves lower banks' funding costs, potentially stimulating credit supply - though this channel is debated.

Quantitative Easing: Mechanisms and Assessment

- **Is QE inflationary?** Not necessarily - and this is where our models matter.
- Under the Quantity Theory, a permanent expansion of the money supply is inflationary. But if QE is expected to be reversed (temporary reserves), and fiscal policy is passive (surpluses adjust), other models (that we will study) suggest the inflationary impact is limited.
- The post-2021 inflation surge reignited this debate: most evidence points to the combination of QE *and* large fiscal transfers (not QE alone) as the key driver of inflation - consistent with the fiscal-monetary interaction framework we now turn to.

Outline

1. A Primer on Monetary Economics
2. Foundations of Money
3. The Welfare Cost of Inflation
4. Deeper Foundations of Money
5. Quantity Theory, Expectations, and Policy
6. Time Inconsistency and Central Bank Design
7. Monetary-Fiscal Interactions
8. The Role of Private Banks

The Government Budget Constraint

- The starting point for any analysis of monetary-fiscal interactions is the **consolidated government budget constraint**. In each period, total government expenditures must be financed by one of three sources: taxes, new borrowing, or money creation. In real terms:

$$\underbrace{G_t - T_t}_{\text{primary deficit}} + \underbrace{(1 + r_{t-1})\, b_{t-1}}_{\text{debt service}} = \underbrace{\Delta b_t}_{\text{new borrowing}} + \underbrace{\sigma_t}_{\text{seigniorage}}$$

  where $b_t = B_t/P_t$ is the real stock of government bonds, r_{t-1} is the real interest rate, and $\sigma_t = (M_t - M_{t-1})/P_t$ is real seigniorage.

- Iterating this constraint forward and imposing the **no-Ponzi condition** (debt cannot grow forever faster than the interest rate), we obtain the **intertemporal budget constraint**:

$$b_0 = \sum_{t=1}^{\infty} \frac{(T_t - G_t) + \sigma_t}{(1 + r)^t}$$

- This equation is the foundation of everything in this section. It says: **any fiscal gap must eventually be closed**. The question is not *whether* it will be closed, but *how* - through future spending cuts, tax increases, or money creation - and *when*.

Seigniorage: Revenue from Money Creation

- **Seigniorage** is the real revenue the government earns by issuing money. In continuous time, with $\mu = \dot{M}/M$ (money growth rate) and $m = M/P$ (real money balances):

$$\sigma = \frac{\dot{M}}{P} = \frac{\dot{M}}{M} \cdot \frac{M}{P} = \mu \cdot m$$

  Seigniorage equals the **money growth rate** (the "tax rate") times **real money demand** (the "tax base").

- This is why seigniorage is called the **inflation tax**. When the government prints money, prices rise, and the real value of money balances held by the private sector falls. The government gains purchasing power; private money holders lose it. The loss is proportional to nominal money holdings - it is a tax on liquidity.

- Equivalently, in steady state where $\pi = \mu$, seigniorage can be written as:

$$\sigma = \pi \cdot m(\pi)$$

  Inflation π is the tax rate; $m(\pi)$ is the tax base, which **shrinks as inflation rises** - rational agents economize on cash holdings to avoid the tax. This immediately suggests a Laffer curve.

The Inflation Laffer Curve

- Seigniorage revenue $\sigma(\mu) = \mu \cdot m(\mu)$ is the product of two forces moving in opposite directions: the tax rate μ rises with money growth, while the tax base $m(\mu)$ shrinks as agents flee inflation. This tension produces an **inflation Laffer curve**.

- Using the Cagan (1956) money demand specification $m(\pi) = Ae^{-\eta\pi}$ (with $\pi = \mu$ in steady state), seigniorage is:

$$\sigma(\mu) = A\mu\, e^{-\eta\mu}$$

Differentiating: $\sigma'(\mu) = Ae^{-\eta\mu}(1 - \eta\mu)$. Revenue is maximized at:

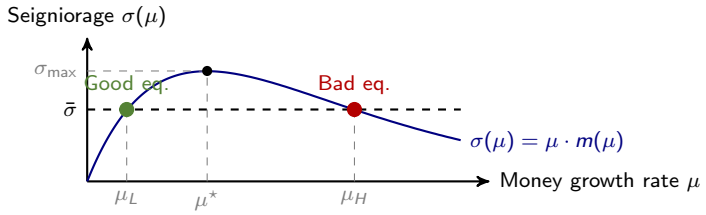
$$\mu^\star = \frac{1}{\eta}$$

For $\mu < \mu^\star$, more money growth raises revenue (tax rate effect dominates). For $\mu > \mu^\star$, more money growth *reduces* revenue (base erosion dominates).

- A government requiring seigniorage $\bar{\sigma} < \sigma(\mu^\star)$ faces **two solutions**: a low-inflation equilibrium $\mu_L < \mu^\star$ (on the "good" side) and a high-inflation equilibrium $\mu_H > \mu^\star$ (on the "bad" side). Both generate the same revenue. The high-inflation equilibrium is **dynamically unstable**: a small shock to expectations triggers a spiral toward hyperinflation.

The Inflation Laffer Curve: A Diagram

- The figure below illustrates the inflation Laffer curve. Seigniorage $\sigma(\mu) = \mu \cdot m(\mu)$ rises then falls with the money growth rate μ. The peak is at $\mu^\star = 1/\eta$. For a required revenue level $\bar{\sigma}$, two equilibria exist: a **good (low-inflation)** equilibrium at μ_L and a **bad (high-inflation)** equilibrium at μ_H.



- The bad equilibrium μ_H is **dynamically unstable**: if agents expect higher inflation, they reduce money demand, which erodes the tax base and forces even more money creation to meet $\bar{\sigma}$ - a self-reinforcing spiral toward hyperinflation.

The Cagan Model of Hyperinflation

- Cagan (1956) studied seven hyperinflations (Germany, Austria, Hungary, Poland, Russia, Greece, China) and proposed a model with one key equation - the money demand function:

$$\ln\left(\frac{M_t}{P_t}\right) = \alpha - \eta\,\pi_t^e, \qquad \eta > 0$$

  Real money demand falls exponentially with expected inflation. When households anticipate rapid price increases, they spend money as fast as possible, reducing real balances.

- With rational expectations ($\pi_t^e = E_t[\pi_{t+1}]$) and a given money supply path, the price level today depends on *the entire expected future path of money growth*. This has a striking implication: **a credible announcement of future fiscal reform can stop inflation immediately** - even before the reform takes effect. Conversely, if the public does not believe that fiscal adjustment is coming, inflation continues even if money growth temporarily slows.

- Sargent (1982) verified this: the German, Austrian, Hungarian, and Polish hyperinflations of the 1920s ended abruptly - within weeks - not through a gradual slowdown of money creation but through **credible fiscal reform**.

- The lesson: **hyperinflation is fundamentally a fiscal phenomenon**.

Sargent-Wallace (1981): The Setting

- The conventional view in the early 1980s was straightforward: if you want low inflation, keep money growth low. The central bank controls inflation; fiscal policy is a separate matter. Sargent and Wallace (1981) showed this view is dangerously incomplete in an economy where the **fiscal authority dominates**.

- Their model has two authorities with conflicting priorities:

  - **The fiscal authority** sets real government expenditures and tax revenues independently, generating a fixed primary deficit $d > 0$ that it will not adjust. It is the **Stackelberg leader**: it moves first and forces the central bank to accommodate.
  - **The central bank** must finance whatever residual the fiscal authority leaves. Its choice is not *whether* to monetize the deficit, but *when* - now or later.

Sargent-Wallace (1981): The Setting

- One additional ingredient drives the result: the **real interest rate on government debt exceeds the economy's growth rate**, $r > g$. This means that unpaid debt grows faster than the economy - the debt-to-GDP ratio is explosive if no primary surplus or seigniorage is forthcoming. The government cannot roll over its debt forever; at some point, it must be paid for.

- In this setting, the central bank faces not a choice between inflation and price stability, but a choice between **inflation now** and **even more inflation later**.

Sargent-Wallace: Unpleasant Monetarist Arithmetic

- Suppose the central bank decides to be **tight today**: it keeps money growth low, generating little seigniorage σ_1 in period 1. The government's deficit must still be financed, so more bonds are issued. The real debt b_1 grows.

- Since $r > g$, this larger debt **compounds over time**. The intertemporal budget constraint must still hold:

$$b_0 = \sum_{t=1}^{\infty} \frac{(T_t - G_t) + \sigma_t}{(1+r)^t}$$

  If the fiscal authority refuses to raise surpluses $(T_t - G_t)$, then the central bank must provide *more* seigniorage σ_t in the future to compensate for the larger stock of debt accumulated today.

Sargent-Wallace: Unpleasant Monetarist Arithmetic

- The arithmetic: **tight money today means loose money tomorrow - and because of the interest compounding, tomorrow's required seigniorage is even larger than what would have been needed today**. A policy of monetary restraint, if not accompanied by fiscal adjustment, merely postpones inflation and makes it worse.

- This is the "unpleasant" result: **in a fiscally dominant regime, open market operations can be inflationary**. Selling bonds to reduce the money supply today simply adds to the debt burden, increasing the future monetary expansion needed to service it. The present value of inflation is unchanged; only its timing shifts - and unfavorably so.

Sargent-Wallace: The Policy Lesson

- The Sargent-Wallace result does *not* say that tight monetary policy is always counterproductive. It says that monetary policy **cannot permanently control inflation when fiscal policy is unsustainable**. The central bank can choose the timing of inflation, but not its present value - that is determined by the fiscal deficit.

- The conditions under which the result holds are important:
  - The fiscal authority must be a true Stackelberg leader: it will not adjust deficits in response to monetary policy.
  - The real interest rate on debt must exceed growth: $r > g$. If $r < g$, the government can roll over debt indefinitely and the problem disappears.
  - There must be an upper bound on seigniorage (the Laffer curve): you cannot collect infinite inflation tax revenue.

- **The practical lesson for central banks**: monetary policy independence is necessary but not sufficient for price stability. A central bank can be formally independent and still fail to control inflation if the fiscal situation is unsustainable - because markets will anticipate future monetization and inflation expectations will rise today.

Fiscal Theory of the Price Level: Motivation

- In the Quantity Theory, the price level is determined by $MV = PY$: money supply pins down prices, and fiscal policy is irrelevant as long as the government eventually balances its budget. In Sargent-Wallace, fiscal policy matters because it forces future money creation. But there is a third, more radical view: **the price level is determined by fiscal policy directly**, regardless of monetary policy.

- This is the **Fiscal Theory of the Price Level (FTPL)**, developed by Leeper (1991), Sims (1994), Woodford (1994, 1995), and Cochrane (2001). The FTPL inverts the usual logic: instead of fiscal policy adjusting to ensure debt sustainability at whatever price level monetary policy determines, *the price level adjusts to ensure that fiscal policy is sustainable at given paths of surpluses and debt*.

Fiscal Theory of the Price Level: Motivation

- The distinction hinges on a simple but profound question: **who adjusts when the government's budget doesn't balance?**
  - **Traditional view (Ricardian)**: fiscal policy adjusts - future surpluses or money creation accommodate whatever is needed. The government's budget constraint holds for any price level; it places no restrictions on P_t.
  - **FTPL (non-Ricardian)**: fiscal policy is fixed - the path of surpluses is set independently and will not change. The price level P_t must then jump to make the constraint hold. The budget constraint becomes a *valuation equation* for the price level.

Leeper (1991): Active and Passive Policies

- Leeper (1991) characterizes monetary and fiscal policy each as either **active** or **passive**, depending on how strongly they respond to the state of the economy.

- **Monetary policy** is summarized by a rule for the nominal interest rate: $i_t = \phi \pi_t + \dots$
  - **Active monetary policy**: $\phi > 1/\beta$ - the real interest rate rises when inflation rises (the Taylor principle). Monetary policy is aggressive enough to stabilize inflation on its own.
  - **Passive monetary policy**: $\phi < 1/\beta$ - the real interest rate does not respond strongly enough to inflation. Monetary policy accommodates whatever the fiscal situation requires.

- **Fiscal policy** is summarized by how taxes respond to debt: $\tau_t = \delta\, b_{t-1} + \dots$
  - **Passive fiscal policy**: δ is large - when debt rises, taxes rise enough to stabilize the debt-to-GDP ratio. Fiscal policy is self-correcting and Ricardian: it ensures solvency for any price level.
  - **Active fiscal policy**: δ is small - taxes do not respond strongly to debt. Fiscal policy sets surpluses independently. Solvency requires the price level to adjust.

Leeper (1991): Active and Passive Policies

- The key insight of Leeper: a well-determined equilibrium requires that **exactly one of the two authorities is active and the other is passive**. When both are active or both are passive, the equilibrium is either overdetermined or indeterminate.
- The four possible combinations of active and passive policy produce strikingly different outcomes:

	Passive Fiscal	**Active Fiscal**
Active Monetary	Unique equilibrium (Standard NK, QTM)	No equilibrium (Overdetermined)
Passive Monetary	Multiple equilibria (Indeterminate)	Unique equilibrium (FTPL regime)

Leeper (1991): The Four Regime Combinations

- **Active monetary + Passive fiscal**: the standard case. Monetary policy controls inflation via the Taylor principle; fiscal policy adjusts taxes to keep debt sustainable. Price level determined by money. This is the regime most central banks aim for.

- **Passive monetary + Active fiscal**: the FTPL regime. Fiscal policy sets surpluses independently; monetary policy accommodates. The price level is determined by the government valuation equation - fiscal fundamentals, not money. This can generate inflation even without money growth.

- The **off-diagonal cells** are pathological: when both authorities try to control the economy independently (both active), there is generically no solution. When both are passive (both accommodate), equilibrium is indeterminate. Good institutional design means **coordinating which authority leads**.

Woodford's Government Valuation Equation

- The centerpiece of the FTPL is the **government valuation equation**, derived from iterating the government's budget constraint forward and imposing the transversality condition. In real terms:

$$\frac{B_{t-1}}{P_t} = \mathbb{E}_t \sum_{s=0}^{\infty} \frac{T_{t+s} - G_{t+s}}{(1+r)^s}$$

The **real value of outstanding nominal debt** must equal the **present value of future primary surpluses**.

- In the traditional (Ricardian) view, this equation always holds because fiscal policy adjusts - it is an equilibrium condition, not a restriction on P_t. In the **FTPL**, the path of surpluses $\{T_{t+s} - G_{t+s}\}$ is fixed exogenously by fiscal policy. The left-hand side contains the unknown: P_t. So the equation **determines** the price level:

$$P_t = \frac{B_{t-1}}{\mathbb{E}_t \sum_{s=0}^{\infty} \frac{T_{t+s} - G_{t+s}}{(1+r)^s}}$$

Higher nominal debt B_{t-1} or lower expected surpluses $\rightarrow$ higher P_t (inflation). Lower debt or fiscal tightening $\rightarrow$ lower P_t (deflation or disinflation).

FTPL vs. the Quantity Theory: The Key Distinction

- The FTPL and the Quantity Theory both explain the price level, but through entirely different mechanisms. Understanding the distinction is essential for policy:

	Quantity Theory	**FTPL**
Price level determined by	Money supply M	PV of fiscal surpluses
Role of fiscal policy	Irrelevant (Ricardian)	Central (non-Ricardian)
Role of money	Essential	Secondary
Inflation without $\uparrow M$?	No	Yes
QE inflationary?	Yes (if permanent)	Not necessarily
Austerity disinflationary?	No	Yes

FTPL vs. the Quantity Theory: The Key Distinction

- **A concrete example**: suppose the central bank conducts quantitative easing (buys bonds, expands money supply) but the fiscal authority simultaneously commits to large future surpluses to retire the debt. Under the QTM, QE is inflationary (more M). Under the FTPL, the price level may not move - the PV of surpluses is unchanged or even increased, so the valuation equation is satisfied at the same P_t.

- **Which theory is right?** The empirical evidence is mixed. The FTPL has gained attention for explaining episodes where inflation appeared disconnected from money growth - including some post-2020 inflation dynamics. Most macroeconomists view the two theories as complementary: the relevant one depends on the *fiscal regime* (Leeper's active vs. passive classification).

Exchange Rate Regimes and Fiscal Constraints

- The choice of exchange rate regime is not purely a monetary decision - it has profound fiscal implications. A **fixed exchange rate** (peg, currency board, or monetary union) imposes a fiscal constraint that a flexible rate does not: it eliminates the inflation tax as a financing option.

- Under a fixed exchange rate and purchasing power parity ($P_t = e\,P_t^\star$, with e fixed and foreign prices $P_t^\star$ stable), the domestic price level is anchored from abroad. The government cannot print money to raise prices, because the resulting inflation differential would make domestic goods uncompetitive, generating balance-of-payments pressure and reserve losses.

- The exchange rate regime thus determines **who disciplines whom**:
  - **Flexible rate**: the central bank can choose any inflation rate; fiscal policy must ultimately be consistent with that choice. Discipline (if any) comes from the central bank.
  - **Fixed rate**: the exchange rate anchors the price level. Discipline is imposed externally - but if the fiscal deficit is large enough, the peg becomes *unsustainable* and eventually collapses.

- The fragility of fixed exchange rates under fiscal pressure is the subject of the classic first-generation currency crisis model. A country that tries to use a fixed rate as a nominal anchor while running persistent deficits is building toward a crisis - the only question is the timing.

Krugman (1979): The First-Generation Crisis Model

- Krugman (1979) provides the classic model of a balance-of-payments crisis driven by fiscal fundamentals. The setup:
  - The government runs a **persistent primary deficit** $d > 0$ that it finances by printing domestic money.
  - The central bank maintains a **fixed exchange rate** $\bar{e}$ by selling foreign reserves whenever the domestic currency faces depreciation pressure.
  - The central bank has a **finite stock of reserves** R_0. When reserves reach zero, the peg must be abandoned.
- The fundamental incompatibility: money creation (needed to finance the deficit) raises domestic prices. To maintain PPP at the fixed rate ($P_t = \bar{e}\, P^\star$), the central bank must sell reserves - absorbing the excess money supply. But reserves are finite. The peg is therefore **unsustainable in the long run**: the deficit is the ticking clock.
- How does the crisis actually unfold? Under rational expectations, the collapse does **not** happen when reserves naturally run to zero. It happens earlier, through a **speculative attack**: agents anticipate the collapse and race to sell domestic currency before the peg breaks, exhausting reserves all at once. The crisis is endogenous to expectations.

Krugman (1979): The Speculative Attack

- To find the attack time, define the **shadow floating exchange rate** $\tilde{e}_t$: the exchange rate that would prevail if the peg were abandoned *right now* and the currency were allowed to float. Under a floating rate, domestic money growth μ (dictated by the fiscal deficit) would cause the exchange rate to depreciate at the same rate: $\tilde{e}_t = \bar{e}_0 \cdot e^{\mu t}$.
- At the moment the peg is maintained, $e_t = \bar{e}$. But $\tilde{e}_t$ is rising over time as the money printing continues. The **speculative attack occurs at the unique time** $T^\star$ **when** $\tilde{e}_{T^\star} = \bar{e}$:
  - **Before** $T^\star$: $\tilde{e}_t < \bar{e}$ - attacking the peg would require selling domestic currency *below* the post-float rate. No rational investor attacks.
  - **At** $T^\star$: $\tilde{e}_{T^\star} = \bar{e}$ - abandoning the peg causes no discontinuous jump in the exchange rate. Investors are indifferent; a coordinated attack is possible with no capital loss. The attack happens; remaining reserves are exhausted instantaneously.
  - **After** $T^\star$: the exchange rate floats and depreciates at rate μ.
- **Key result**: the crisis is *triggered* by rational agents, but *caused* by the fiscal deficit. The government cannot prevent the crisis by simply defending the peg more aggressively - it can only delay it by draining reserves faster. The only genuine fix is **fiscal adjustment**: eliminating the deficit that makes money creation necessary.

Dollarization: Commitment by Abandonment

- The first-generation crisis model shows that a fixed exchange rate is fragile when fiscal policy is unsustainable. The most extreme solution to this fragility is **dollarization**: *abandon the domestic currency entirely* and adopt a foreign currency (typically the US dollar) as legal tender. Ecuador (2000), El Salvador (2001), and Panama (since 1904) are the leading examples.

- **The gains from dollarization**:
  - **Eliminates exchange rate risk**: firms and households can borrow and contract in the currency they transact in. Risk premia fall; interest rates converge toward US rates.
  - **Hard commitment on inflation**: the government literally cannot print money to finance deficits. The inflation tax disappears. This can be highly credible when domestic institutions are weak.

- **The costs of dollarization**:
  - **Loss of seigniorage**: the government foregoes the revenue from money creation, which for some developing countries is a significant share of fiscal revenues.
  - **No lender of last resort**: the central bank cannot create dollars to rescue failing banks. Financial crises become harder to manage.

Exchange Rates and Prices in Mexico

- Mexico's Central Bank has published several papers that estimate the relationship between prices and exchange rates (mostly with respect to the US dollar).
- Some insights on this relationship:
  1. There was a one-to-one co-movement between the US-dollar and the peso up to the early 2000s.
     - ⋆ Diaz de León et al. (2003), Lopez-Martín et al. (2018).
  2. Between 2000-2020 there seems to be a somewhat disconnection between prices and exchange rates.
     - ⋆ Estimates show a pass-through between 0.04 - 0.15 percentage points for every 1% depreciation of the peso.
     - ⋆ Capistrán et al. (2011), Cortés (2013), Kochen and Sámano (2016).
     - ⋆ These estimates are quite low!
- How can we rationalize this behavior?

Exchange Rate and Prices in Mexico

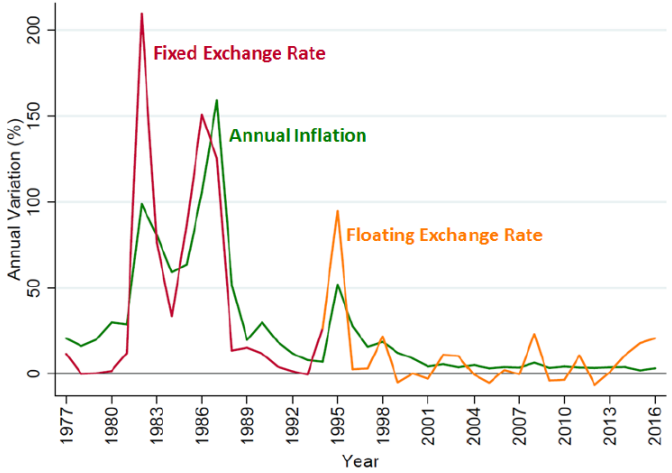


Figure: Inflation and Exchange Rate Dynamics in Mexico.

Outline

1. A Primer on Monetary Economics
2. Foundations of Money
3. The Welfare Cost of Inflation
4. Deeper Foundations of Money
5. Quantity Theory, Expectations, and Policy
6. Time Inconsistency and Central Bank Design
7. Monetary-Fiscal Interactions
8. The Role of Private Banks

Banks and the Macroeconomy: Why They Matter

- All the models we have studied so far feature a central bank and a representative household - but no private banks. This is a significant omission. Private banks are not merely conduits for central bank policy: they are **active creators of money, credit, and liquidity**, and their health is inseparable from price stability.

- **Banks create inside money.** When a bank makes a loan, it credits the borrower's deposit account - creating a new deposit without any prior saving. The broad money supply ($M2$, $M3$) is therefore largely determined by private bank balance sheets, not by the central bank alone. The money multiplier connects the monetary base to the total money supply: a contraction in bank credit shrinks the money supply independently of central bank actions.

Banks and the Macroeconomy: Why They Matter

- **The bank lending channel.** Monetary policy affects the economy partly through its effect on bank credit supply. When the central bank raises rates, bank funding costs rise, credit standards tighten, and investment falls - even for borrowers who are not directly rate-sensitive. Bernanke (1983) showed that the collapse of the US banking system in 1930–1933 *deepened and prolonged* the Great Depression far beyond what interest rates alone can explain.

- **Financial stability as a prerequisite for price stability.** A banking crisis destroys the monetary transmission mechanism: insolvent banks do not pass on rate cuts; hoarding of reserves prevents money creation. This is why central banks function as **lenders of last resort** and why banking regulation is inseparable from monetary policy.

Diamond (1984): The Delegated Monitoring Problem

- Diamond (1984, *Review of Economic Studies*) asks a foundational question: **why do banks exist?** The answer lies in an information problem between borrowers and lenders.

- **The environment**: an entrepreneur needs \$1 to fund a project. The project yields a random return R_i, but the outcome is *private information* - only the entrepreneur can observe it without a costly audit. Each lender who wants to verify the return must pay a monitoring cost $c > 0$.

- **The problem with direct lending**: suppose n investors each lend $1/n$ to the entrepreneur. Each investor who wants to enforce repayment must monitor independently, at cost c. The **total monitoring cost** is nc. For large n, this is prohibitively expensive - a massive duplication of effort. Each investor free-rides on others' monitoring, and in equilibrium, no one monitors. The entrepreneur has no incentive to repay.

- The problem is a classic **moral hazard with costly state verification**: lenders cannot commit to auditing cheaply, and borrowers exploit this. Direct lending to anonymous entrepreneurs is inefficient in large economies. **Something must reduce or eliminate the duplication of monitoring costs** - and that something is the bank.

Diamond (1984): Banks as Delegated Monitors

- **The solution**: a bank borrows $1/n$ from each of the n investors, pools the funds, lends \$1 to the entrepreneur, and monitors *once* on behalf of all investors. The total monitoring cost is now simply c - independent of n. The cost per investor falls to c/n, which vanishes as the bank grows.

- But this raises a new question: **who monitors the monitor?** If investors delegate monitoring to the bank, what prevents the bank from misreporting the entrepreneur's returns? Diamond's insight is that this problem is solved by **diversification**. A bank that lends to many independent entrepreneurs has a portfolio return that is nearly deterministic by the law of large numbers:

$$\bar{R} = \frac{1}{K} \sum_{k=1}^{K} R_k \xrightarrow{K \to \infty} \mathbb{E}[R_k]$$

  With a large, diversified portfolio, the bank can offer depositors a **fixed-return debt contract** (a deposit): "pay me d regardless of the loan portfolio outcome." This debt contract is *incentive-compatible* for the bank - it has no reason to misreport since its promise to depositors is fixed.

Diamond (1984): Banks as Delegated Monitors

- **The result**: banks emerge endogenously as the efficient solution to the delegated monitoring problem. They hold *opaque*, *risky*, *illiquid* assets (loans to entrepreneurs) and issue *transparent*, *safe*, *liquid* liabilities (deposits). This is the essence of banking - and it is also the source of their fragility, as Diamond and Dybvig (1983) show.

Banks as Secret Keepers: Gorton, Holmström, and Ordoñez (2017)

- Gorton, Holmström, and Ordoñez (2017) argue that the opacity of bank assets is not merely a feature of the financial system - it is **socially optimal**. Banks are, in their words, **"secret keepers."**

- **The key concept**: an asset is **informationally insensitive** when the gains from investigating its quality are smaller than the cost of doing so. In this case, agents accept the asset *without due diligence* - on a "no-questions-asked" basis. Bank deposits and short-term bank paper function this way: people accept them as payment without auditing the bank's loan portfolio.

- **The argument for opacity**: suppose bank assets were fully transparent. If a bank's loan portfolio deteriorates, informed investors would demand a premium or refuse to roll over short-term funding. But information production is costly - and if one agent investigates, others can free-ride. In equilibrium with full transparency, **too many resources are spent on information production**, and the informational insensitivity of bank liabilities is destroyed. Trade in safe assets collapses.

- By keeping assets opaque, banks sustain a **no-investigation equilibrium** in which their liabilities are treated as safe. This is efficient as long as the assets are genuinely safe enough: the expected cost of opacity (occasionally holding bad assets undetected) is smaller than the benefit (preserving the liquid, no-questions-asked property of bank money).

GHO (2017): Fragility, Crises, and Implications

- The GHO framework implies that the financial system is **inherently fragile** - not despite bank opacity, but because of it. When the quality of bank assets deteriorates sufficiently, the no-investigation equilibrium can break down:
  - As assets become riskier, the potential gain from investigating their quality rises. At some threshold, it becomes individually rational for agents to start auditing bank assets.
  - Once investigation begins, bank liabilities become **informationally sensitive**: investors demand compensation for the risk of holding them. The "no-questions-asked" property is destroyed.
  - This is a **discontinuous transition** - a small deterioration in asset quality can trigger a sudden collapse of the entire safe-asset system.

GHO (2017): Fragility, Crises, and Implications

- **The 2008 financial crisis as a GHO event.** Mortgage-backed securities were initially treated as informationally insensitive (rated AAA, accepted as collateral without scrutiny). When housing prices fell and MBS quality became uncertain, investors began investigating. Money market funds "broke the buck," repo markets froze, and interbank lending collapsed - precisely the information cascade GHO describe.

- **Policy implications**: regulators face a dilemma. Forcing banks to be more transparent may seem prudent in normal times, but it destroys the informational insensitivity that makes the financial system function. Stress tests, capital requirements, and deposit insurance work precisely because they *restore* confidence in bank liabilities without requiring investors to investigate underlying assets.

Diamond-Dybvig (1983): The Environment

- Diamond and Dybvig (1983, *Journal of Political Economy*) provide the canonical model of bank runs. Their central insight is that demand deposit contracts - the very mechanism that makes banks useful - also makes them fragile.

- **Environment**: there are three periods ($t = 0, 1, 2$) and a continuum of agents, each endowed with 1 unit of the good at $t = 0$. At $t = 1$, each agent privately learns her **type**:
  - **Impatient** (fraction λ): only values consumption at $t = 1$. Utility: $u(c_1)$.
  - **Patient** (fraction $1 - \lambda$): values consumption at $t = 1$ or $t = 2$. Prefers to consume later if possible.

- **Technology**: there is one investment opportunity. One unit invested at $t = 0$ returns $R > 1$ units at $t = 2$ (a long, illiquid investment). If the investment is **liquidated early** at $t = 1$, it returns only 1 unit (the opportunity cost of illiquidity is the foregone return $R - 1$).

- **The maturity mismatch**: this technology is well-suited for patient agents but costly for impatient ones. In autarky (no bank), impatient agents receive 1 unit and patient agents receive R. The question is whether a bank can do better by **pooling the liquidity risk** across agents whose types are unknown ex ante but independent across individuals.

Diamond-Dybvig: The Optimal Deposit Contract

- A social planner who can observe types would solve:

$$\max_{c_1,\, c_2} \; \lambda\, u(c_1) + (1-\lambda)\, u(c_2) \quad \text{s.t.} \quad \lambda c_1 + \frac{(1-\lambda)c_2}{R} = 1$$

  The optimal solution satisfies $u'(c_1) = R\, u'(c_2)$, which - under risk aversion - implies $1 < c_1^\star < c_2^\star < R$. Impatient agents receive *more* than autarky (1), and patient agents receive *less* than autarky (R). The bank provides **insurance against liquidity risk**.

- A bank implements this with a **demand deposit contract**: any agent may withdraw $c_1^\star$ at $t = 1$ on demand. Agents who wait until $t = 2$ share the remaining assets and receive $c_2^\star > c_1^\star$.

- **The good equilibrium**: if only impatient agents withdraw at $t = 1$, the bank honors all withdrawals and the first-best allocation is achieved. Patient agents wait, the long investment matures, and everyone is better off than in autarky.

- **The problem**: the demand deposit contract also admits a second, catastrophic equilibrium. Because withdrawals are served *sequentially* and on a first-come-first-served basis, a patient agent who believes others will run has a strong incentive to run too - before the bank's assets are exhausted. The same contract that provides liquidity insurance also creates fragility.

Diamond-Dybvig: Bank Runs as Self-Fulfilling Equilibria

- **The bank run equilibrium**: suppose a patient agent believes that *all other* patient agents will withdraw at $t = 1$. Then the bank will serve withdrawals sequentially until its assets are exhausted. A patient agent who waits until $t = 2$ receives nothing (the bank has already paid out $c_1^\star$ to everyone ahead of her in the queue). Withdrawing at $t = 1$ yields $c_1^\star > 0$. Therefore:

$$\text{Run if others run:} \quad c_1^\star > 0 = \text{payoff from waiting}$$

Running is individually rational whenever others run. This is a Nash equilibrium in which all agents - regardless of type - withdraw at $t = 1$. The bank liquidates all its long investments early, destroying the value $R - 1 > 0$ that would have been created by waiting.

Diamond-Dybvig: Bank Runs as Self-Fulfilling Equilibria

- **The key insight**: the bank run is not caused by insolvency. The bank's assets are perfectly sound - if everyone waited, all obligations could be met. The run is a **coordination failure**: a self-fulfilling panic driven entirely by beliefs, not fundamentals. Both equilibria (no run and run) are consistent with the same model, the same bank, and the same asset quality.

- **Sunspot equilibria**: Diamond and Dybvig show that any publicly observable, payoff-irrelevant event (a "sunspot" - bad news, a rumor, a tweet) can coordinate agents onto the run equilibrium. Banks are therefore **inherently vulnerable to self-fulfilling crises**, even when they are fundamentally solvent.

Diamond-Dybvig: Policy Responses

- Diamond and Dybvig analyze three policy responses to bank runs, each with different properties:
  - **Suspension of convertibility**: the bank announces it will stop honoring withdrawals once a fraction λ of deposits have been paid out. Patient agents know they will receive $c_2^\star$ by waiting, so running is no longer rational. The run equilibrium is *eliminated*. However, suspension is socially costly if some genuinely impatient agents arrive after the threshold - they cannot access their funds. It also requires the bank to know λ in advance.
  - **Deposit insurance**: the government guarantees that every depositor will receive $c_1^\star$ at $t = 1$ regardless of the bank's condition. A patient agent now has no incentive to run: she will receive $c_2^\star > c_1^\star$ by waiting, and if the bank fails the government makes her whole. The run equilibrium is *eliminated* and the first-best is achieved - as long as the government's guarantee is credible.
  - **Lender of last resort (Bagehot rule)**: the central bank stands ready to lend unlimited amounts to illiquid but *solvent* banks at a penalty rate. Patient depositors know the bank can always honor withdrawals, so running is never rational. This is the classic central bank function - and it connects the Diamond-Dybvig model directly to the institutional design of modern central banks.