LECTURE 1.5

Banks, Bank Runs and the Financial Cycle

ECON42550 Macroeconomics

Sam Deegan
University College Dublin

Autumn 2026

Where We Are

1. Short-Run Fluctuations

1.1 The Keynesian Multiplier

1.2 Introducing IS-MP-PC

1.3 Analysing IS-MP-PC

1.4 The Zero Lower Bound

1.5 Banks, Runs, the Cycle

2. NK Building Blocks

2.1 Expectations, Investment

2.2 Consumption, Fiscal Policy

2.3 Sticky Prices and the NKPC

3. Long-Run Growth

3.1 Growth Facts, Solow

3.2 The Fundamental Causes

3.3 Endogenous Growth

3.4 Schumpeterian Growth

Last Week's Lecture

Last week the policy interest rate ran into a floor:

1. **Why There Is a Bound.**
   Establish that the policy rate cannot fall far below zero, and why that matters.
2. **The Trap.**
   Show what the model does once the bound binds, and why it is self-reinforcing.
3. **Policy at the Bound.**
   Work through what a central bank can still do when the rate is stuck.

Lecture Plan

1. **Why Banks Are Fragile.**
   Why a bank with good assets can still fail.

2. **Runs.**
   Show that the very contract that makes a bank useful has a second, bad equilibrium — and go through the policies that rule it out.

3. **The Financial Cycle.**
   Step back from one bank to the whole system: credit, leverage, and the difference between running out of cash and running out of capital.

> **Diamond, Plus Romer Chapter 10**
> Reading: Diamond's Richmond Fed exposition (Diamond, 2007), plus Romer, chapter 10, both on Brightspace.

SECTION 1

Banks in the Model

Why Banks Matter for Macro

The model so far has one interest rate; this section adds a second. Between the central bank and the firm sits a bank, and the bank charges more:

- Banks fund themselves overnight, at an interest rate the central bank sets.
- A five-year plant loan is riskier and longer, so it costs more.
- The gap widens by percentage points while the policy interest rate stands still.

> **Policy Runs Through a Rate the Bank Does Not Set**
> Everything monetary policy does, it does through an interest rate the central bank does not set.

What a Bank Charges, and Why

Four things sit inside the gap, and the central bank sets none of them:

1. **Funding Cost.** Deposits and wholesale funding do not track the policy interest rate.
2. **Default Risk.** Every interest rate prices in the loans that are never repaid.
3. **Capital and Regulation.** Equity is expensive, and riskier loans require more of it.
4. **Market Power.** A concentrated market makes a loan's price a decision, not a cost.

> **The First Three Widen Together in a Recession**
> A spread that never widened could be folded into r^* and forgotten. This one widens exactly when demand is weakest.

The Mark-Up Breaks the Transmission Mechanism

Firms and households do not borrow at the policy interest rate. They borrow at

$$\ell_t = i_t + \underbrace{\mu_t^B}_{\text{the bank mark-up}} \tag{1}$$

- ℓ_t: the interest rate firms and households borrow at
- μ_t^B: the mark-up banks charge over the policy interest rate

> **Demand Answers to ℓ_t, the Policy Rule to i_t**
> Spending falls with what a borrower pays, not with what a central bank announces. So ℓ_t replaces i_t inside the IS curve from here on.

The Four-Equation Model

Demand, supply, policy, and now intermediation, in that order:

$$y_t = y_t^* - \alpha \underbrace{(\ell_t - \pi_t - r^*)}_{\text{the real borrowing rate gap}} + \varepsilon_t^y \tag{2}$$

$$\pi_t = \pi_t^e + \gamma(y_t - y_t^*) + \varepsilon_t^\pi \tag{3}$$

$$i_t = r^* + \pi^* + \beta_\pi(\pi_t - \pi^*) \tag{4}$$

$$\ell_t = i_t + \mu_t^B \tag{5}$$

Four Equations, Four Unknowns

Only the first equation has changed, and μ_t^B is exogenous for now.

Four Equations, and the System Closes

The four unknowns are output, inflation, the policy interest rate and the borrowing interest rate.

Demand, With the Bank In It

The curve now runs on the borrowing rate, and that gap splits in two:

$$y_t = y_t^* - \alpha \left(\ell_t - \pi_t - r^*\right) + \varepsilon_t^y \tag{6}$$

$$\ell_t - \pi_t - r^* = (\beta_\pi - 1)(\pi_t - \pi^*) + \mu_t^B \tag{7}$$

The split goes in where the gap was:

$$y_t = y_t^* - \alpha \left[(\beta_\pi - 1)(\pi_t - \pi^*) + \mu_t^B \right] + \varepsilon_t^y \tag{8}$$

The Mark-Up Sets the Curve's Position

β_π still sets the **slope**; a rise in μ_t^B shifts the curve **left** by α times that rise.

What Would Make a Bank Widen Its Spread?

The model has taken μ_t^B as given. Nothing written down so far says what sets it.

> **Question**
> Of the four things inside the spread, which of them widen when output falls? What does that do to demand, with the spread now inside the IS curve?

Take answers from the room.

The Mark-Up Widens When Output Falls

$$\mu_t^B = \bar{\mu}_t^B - \underbrace{\phi(y_t - y_t^*)}_{\text{the financial accelerator}} \tag{9}$$

- $\bar{\mu}_t^B$: the part of the mark-up set outside the model
- ϕ: how far the mark-up falls when output rises by one point

A Bar Means Given from Outside

$\bar{\mu}_t^B$ is banks' funding conditions, their capital, how frightened they are. A banking crisis, in this model, is a rise in $\bar{\mu}_t^B$.

The Loop Amplifies Whatever Started It

Output falls, borrowers look riskier, the mark-up widens, output falls further. The loop starts no recession; it deepens the one that started.

How Much the Accelerator Amplifies

The curve with the bank in it, and the mark-up that answers to output:

$$y_t - y_t^* = -\alpha(\beta_\pi - 1)(\pi_t - \pi^*) - \alpha\,\mu_t^B + \varepsilon_t^y \tag{10}$$

$$\mu_t^B = \bar{\mu}_t^B - \phi(y_t - y_t^*) \tag{11}$$

The mark-up lands where it stood, bringing a term in y_t with it:

$$y_t - y_t^* = -\alpha(\beta_\pi - 1)(\pi_t - \pi^*) - \alpha\left[\bar{\mu}_t^B - \phi(y_t - y_t^*)\right] + \varepsilon_t^y \tag{12}$$

Collect the output terms and divide by $1 - \alpha\phi$:

$$y_t - y_t^* = \frac{-\alpha(\beta_\pi - 1)(\pi_t - \pi^*) - \alpha\bar{\mu}_t^B + \varepsilon_t^y}{1 - \alpha\phi} \tag{13}$$

Every Shock Is Divided by $1 - \alpha\phi$

What a Wider Spread Does to Output and Inflation

Banks' funding conditions deteriorate and the mark-up $\bar{\mu}_t^B$ rises by two percentage points. The central bank leaves i_t exactly where it was.

> **Question**
> What happens to output and to inflation? Both **up**, both **down**, or one of each?

Think on your own first, then take ninety seconds with the person beside you.

A Mark-Up Shock, Traced

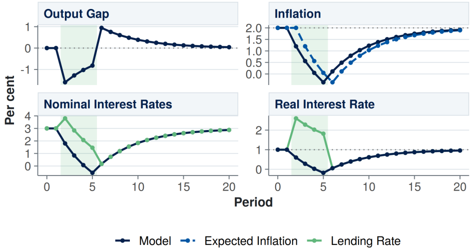

Figure 1: Four time paths after a two-point rise in $\bar{\mu}_t^B$: the lending interest rate jumps above the policy interest rate, the output gap falls, inflation falls, and the policy interest rate is cut in response.

> **No Policy Decision Caused This Contraction**
> Both series fall on impact. Falling output widens the mark-up again through ϕ, so the recovery lags an equivalent demand shock.

What the Central Bank Can Do About It

The rule cuts i_t only once inflation has fallen, and inflation falls last. Three responses, in ascending order of how far they leave the model:

- **Cut Below the Rule:** Offset the rise in μ_t^B one for one with i_t.
- **Lend Directly:** Fund the banks that cannot fund themselves — a quantity, not a price.
- **Regulate the Causes:** Capital rules and supervision hold $\bar{\mu}_t^B$ down, slowly.

> **The Model Gives the Policy, Not the Measurement**
> Cutting below the rule is exact in the model and hard in practice. $\bar{\mu}_t^B$ has to be observed in real time.

Broken Transmission in Europe

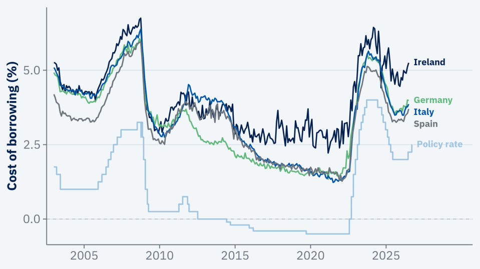

Figure 2: New corporate loan rates, selected euro-area countries, against the ECB policy rate.

> **The Same i_t, Different ℓ_t, Different Demand**
> The flat line is the policy interest rate, common to everyone. Each country's vertical distance above it is that country's μ_t^B.

Two Additions to the Model

Both arrived today, and together they separate the policy rule from borrowing costs:

1. **The Bank Mark-Up.**
   Firms borrow at $\ell_t = i_t + \mu_t^B$, not at the policy interest rate.
2. **The Financial Accelerator.**
   The mark-up widens when output falls, so every shock is divided by $1 - \alpha\phi$.
3. **Together.**
   A central bank can satisfy $\beta_\pi > 1$ and still not control what firms pay.

SECTION 2

Why Banks Are Fragile

What a Bank Actually Does

This section builds the model that explains why a solvent bank can still fail. Strip a bank down and it is doing one thing, over and over:

- A bank takes money from people who want it back **at any moment**.
- The bank puts that money into things that pay off **years later**.

The Mismatch Is the Design

Borrowing short and lending long. The liabilities can be called tomorrow, and the assets cannot be sold tomorrow without a loss.

Usefulness and Danger Share One Feature

Everything useful about a bank and everything dangerous come from the mismatch. No bank keeps the usefulness and sheds the danger.

Liquidity, and Who Provides It

You do not know when you will need your money. A broken car, a flat deposit, a job that falls through — or not for five years:

- Held as **cash**, your money is available whenever you want it and earns nothing.
- Put into a **project**, it earns a return, but only if you leave it there.

You would give up some return for the option of getting out early. That is an insurance contract, and a bank writes it.

> **Liquidity Is Spending Power, Available Now**
> The ability to turn an asset into spending power **now**, at close to its full value. Cash is perfectly liquid. A half-built factory is not.

The Setup: Three Dates, One Asset, Two Exits

Three dates, one unit of goods each at $t = 0$, and one asset with two exits:

$$1 \text{ unit at } t = 0 \longrightarrow \begin{cases} 1 & \text{if taken out at } t = 1 \quad \text{(stopped early)} \\ R > 1 & \text{if left until } t = 2 \quad \text{(the long asset)} \end{cases} \tag{14}$$

- t: the date, 0, 1 or 2
- R: the gross return on the long asset at $t = 2$

Subscripts Are Dates, and Consumption Stays Lower Case

Every past paper writes c_1 for date-1 consumption, in lower case.

Interrupting the Long Asset Destroys $R - 1$

Taken out at $t = 1$ the long asset returns the unit that went in, and $R - 1$ is lost.

The Setup: Two Kinds of Depositor

At $t = 0$ nobody knows what they will need. At $t = 1$ each depositor learns their type:

$$U = \begin{cases} u(c_1) & \text{impatient, with probability } \lambda \\ u(c_2) & \text{patient, with probability } 1 - \lambda \end{cases} \tag{15}$$

- λ: the fraction of depositors who turn out to be impatient
- c_1: consumption at date 1
- c_2: consumption at date 2
- $u(\cdot)$: utility from consumption, increasing and concave

> **Read λ Where a Past Paper Writes π**
> π has meant inflation in this module since Lecture 1.2 and goes on meaning it. The impatient fraction is λ here.

Why the Risk Is Worth Insuring

The long asset punishes the impatient. They recover one unit; the patient collect R. Two facts make that risk insurable:

- Nobody knows at $t = 0$ which of the two they will be.
- A **fraction** λ of the population will be impatient, and λ is known.

Idiosyncratic, Not Aggregate
Each person faces uncertainty; the population faces none. Insurance exploits that gap.

Pooling Beats Acting Alone
With λ known, a bank can pay the impatient more than one unit.

Autarky: What You Get on Your Own

With no bank, each person invests their own unit in the long asset. The impatient pull it out at date 1.

$$c_1 = 1, \qquad c_2 = R \tag{16}$$

The impatient bear the whole cost of the mismatch.

Autarky Is Inefficient, Not Unfair
Expected utility is $\lambda\, u(1) + (1 - \lambda)\, u(R)$. Everyone faced the same lottery, and the population carries no risk.

The Optimal Allocation: The Problem

Put a planner in charge at $t = 0$, before anyone knows their type. The planner maximises expected utility subject to the technology:

$$\max_{c_1,\,c_2} \;\lambda\, u(c_1) + (1-\lambda)\, u(c_2) \qquad \text{s.t.} \qquad \lambda c_1 + \underbrace{\frac{(1-\lambda)\, c_2}{R}}_{\text{today's cost of the late payment}} = 1 \tag{17}$$

Why the Second Term Is Divided by R

A unit left in the long asset grows to R by date 2. Delivering c_2 then costs only c_2/R today.

The Optimal Allocation: The Answer

One more unit paid at date 1 costs R units at date 2:

$$u'(\hat{c}_1) = \underbrace{R\,u'(\hat{c}_2)}_{\text{the marginal value of waiting}} \tag{18}$$

- $\hat{c}_1$: the optimal payment to an impatient depositor
- $\hat{c}_2$: the optimal payment to a patient depositor

A Hat Is an Optimum; a Star Is Still a Resting Point

A star marks where the economy **settles**. A hat marks the solution to a maximisation.

The Planner Takes From the Patient

With relative risk aversion above one, $1 < \hat{c}_1$ and $\hat{c}_2 < R$.

The Bank as Intermediary

A bank promises r_1 at date 1, and the rest at date 2. Paying the impatient costs λr_1, so what stays in the ground is

$$\underbrace{1 - \lambda r_1}_{\text{the assets still invested}} \tag{19}$$

- r_1: what the contract pays a depositor withdrawing at date 1

The contract is **incentive compatible**: each type prefers the payment meant for it.

> **r_1 Is a Payment, Not an Interest Rate**
> r_1 and r_2 are goods a depositor collects at each date, not the r^* of Lecture 1.3. The only price in this model is R, and it is a technology.

The Patient Pay for the Impatient

What stays in the ground grows by R and is shared among the $1-\lambda$ depositors who waited:

$$r_1 = \hat{c}_1, \qquad r_2 = \frac{R\,(1-\lambda r_1)}{1-\lambda} \tag{20}$$

- r_2: what the contract pays a depositor withdrawing at date 2

> **$r_1 > 1$ Beats Autarky, and $r_2 < R$ Pays for It**
> The impatient withdraw early because they want to. The patient wait because $r_2 > r_1$.

The Contract, in Numbers

The module's calibration is the one behind the Diamond–Dybvig app. The sample paper uses the same numbers.

> **The Impatient Get Back 28% More Than They Put In**
> $r_1 = 1.28$ at date 1, where the technology alone returns exactly one unit.

Where the Extra 28% Comes From

The contract pays an early withdrawer $r_1 = 1.28$. The technology on its own would have paid exactly 1.

> **Question**
> Who pays for the extra 0.28, and what do they get in exchange for paying it?

Take answers from the room. Then open the Diamond–Dybvig app at stage 3 and read r_2 off it.

SECTION 3

Runs

The Same Contract, Read Backwards

This section shows that the same contract has a second, bad equilibrium. Ask what a patient depositor should do if they think **everyone else** is about to withdraw:

- The bank pays r_1 in arrival order at date 1, until it runs out.
- It holds one unit per depositor and owes each of them $r_1 = 1.28$.

> **Sequential Service Rewards Running First**
> Depositors are paid in arrival order, not pro rata: the first in the queue collect r_1 and the last collect nothing.

Two Equilibria: Everyone Waits

Suppose every patient depositor expects the others to wait:

- Only the impatient withdraw at date 1: a fraction λ, taking r_1 each.
- The bank liquidates λr_1 and leaves $1 - \lambda r_1$ invested.
- At date 2 the $1 - \lambda$ patient share $R(1 - \lambda r_1)$, so each gets $r_2 > r_1$.

Waiting Is a Best Response to Waiting

The allocation is $\hat{c}_1, \hat{c}_2$ — the planner's optimum, delivered by a contract nobody had to be forced into.

Two Equilibria: Everyone Runs

Now suppose every patient depositor expects the others to withdraw at date 1:

- If you wait, there will be nothing left, so you get **zero**.
- If you join the queue you get r_1 with some probability, which is more than zero.

$$\text{fraction of depositors paid at } t = 1 \;=\; \frac{1}{r_1} \;<\; 1 \tag{21}$$

So everyone joins the queue. The assets were good and would have paid R.

> **The Bank Failed Because It Was Asked To**
> The bank holds one unit per depositor and owes 1.28, so it fails. Running is a best response to everyone else running.

The Optimal Contract Is What Makes the Run Possible

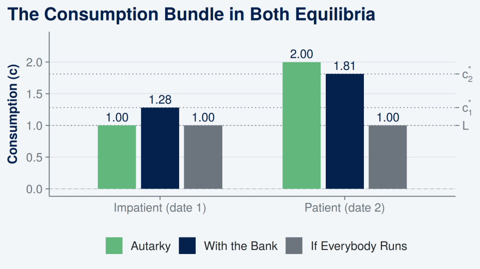

Figure 3: What each depositor type gets under autarky, under the contract, and in a run.

Liquidating everything raises exactly 1 per depositor, so the run equilibrium exists whenever $r_1 > 1$ — and the optimal contract satisfies it, since $r_1 = \hat{c}_1 > 1$ whenever relative risk aversion exceeds one.

Insurance and Fragility Are One Object

The **best** plain deposit is run-prone: run-proof means $r_1 = 1$, which is autarky.

Why Both Are Equilibria

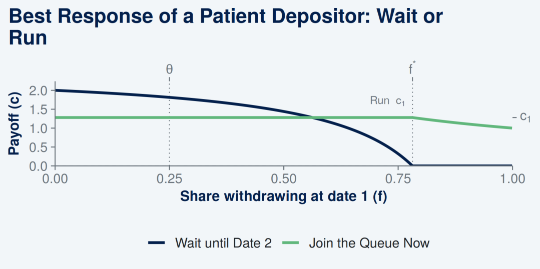

Figure 4: A patient depositor's payoff from waiting, which falls as more depositors withdraw, against the flat payoff c_1 from joining the queue. The two cross well before the point at which the bank fails.

Strategic Complementarity

Each depositor in the queue leaves less for those who wait, so the more who withdraw, the better withdrawing is for you.

Self-Fulfilling

Believing the bank will fail makes it fail.

A Numerical Run, in Three Steps

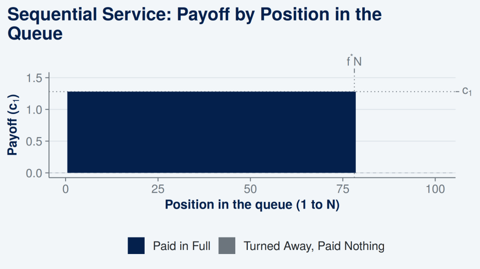

Figure 5: Sequential service: everyone up to f^*N is paid in full; the rest get nothing.

1. **The Contract.**
   Write down r_1 and the date-2 payment r_2.

2. **The Good Equilibrium.**
   A fraction λ withdraws early: show that a patient depositor prefers to wait.

Preventing Runs: Suspension of Convertibility

The contract lets the bank stop paying once a fraction λ has withdrawn:

- A patient depositor may be turned away, while the suspension protects r_2.
- So no patient depositor joins the queue, and the suspension is never used.

The threat works because it is never carried out. Only the good equilibrium is left.

> **Suspension Needs λ to Be Known**
> If the impatient fraction is itself uncertain, suspending at λ turns away people who really did need their money.

Preventing Runs: Deposit Insurance

A government guarantees that every depositor is paid, whatever happens to the bank:

- A patient depositor gets r_2 whether or not anyone else runs.
- The run equilibrium disappears, and again the guarantee is never called on.

The Guarantee Removes the Belief, Not the Risk
Deposit insurance puts no money into the bank. It removes the belief that made the run rational in the first place.

Insured Depositors Stop Watching the Bank

A guarantee takes away the depositor's reason to run. It removes their reason to care what the bank does with the money.

> **Moral Hazard Is Why the Guarantee Is Capped**
> Deposit insurance comes bundled with capital requirements, supervision, and a cap: €100,000 per depositor per institution across the euro area.

Preventing Runs: The Lender of Last Resort

The bank's assets in a run are good; they cannot be turned into cash today. So a central bank lends against them:

- It advances cash against those assets at a **discount**: the gross rate $\hat{c}_2/\hat{c}_1$, below R.
- It lends only beyond the first λ of withdrawals, so the bank cannot borrow to arbitrage the difference.

Bagehot's rule is to lend freely, against good collateral, at a penalty rate. The model agrees on the collateral and contradicts the other two.

> **Three Routes, One Credible Promise**
> All three policies protect the promise to those who **wait**. Nobody then has a reason to be first in the queue.

A Run Spreads Without a Second Mistake

Three routes carry a run from one bank to the next:

1. **Information.**
   A failure elsewhere is evidence about banks in general, and about yours in particular.
2. **Exposure.**
   Banks lend to each other overnight, so a failure is a loss for everyone who lent to it.
3. **Fire Sales.**
   A forced seller marks those assets down everywhere, which forces the next bank to sell.

Fire Sales Reach Banks With No Connection at All

The third route runs through **prices**, and it is what makes one bank's failure a macroeconomic event.

Which Assumption Would Remove the Run?

The run needs three things: sequential service, $r_1 > 1$, and beliefs the model does not pin down.

Question
Drop one of the three. Which one removes the run equilibrium, and what does dropping it cost the depositors?

Take answers from the room before anyone opens a book.

SECTION 4

The Financial Cycle

From One Bank to the Whole System

Diamond–Dybvig (Diamond & Dybvig, 1983) is a model of one bank on one afternoon. The credit cycle behind 2008 ran for a decade:

- Credit grows **faster than output** for years, and then falls faster.
- The build-up is slow and looks benign; the unwinding is fast and disorderly.
- Credit turns **before** output does.

The Financial Cycle
The joint movement of credit, leverage and asset prices. Its swings are measured in decades and are larger than the business cycle's.

Credit and Leverage in the Data

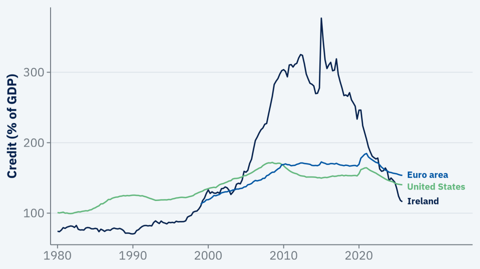

Figure 6: Credit to the private non-financial sector as a share of GDP: Ireland, the euro area and the United States.

> **Credit Peaks Before the Downturn**
> Credit divided by annual output, with recessions shaded. Each credit series turns down **before** its recession band, not during it.

What Absorbs a Bank's Losses

A bank funds its assets mostly by borrowing, and partly with its own money. What is left once everything owed has been paid is

$$\underbrace{A_t - L_t}_{\text{the bank's own capital}} \tag{22}$$

- A_t: the value of the bank's assets
- L_t: what the bank owes to depositors and other lenders

> **Capitals Are Levels, Here Too**
> A_t, L_t and equity are stocks in euros on a date, as Y, C, I and G were in Lecture 1.1.

Leverage, and Why It Amplifies

Leverage counts how many euros of assets each euro of equity supports:

$$\text{leverage} = \frac{A_t}{A_t - L_t} \tag{23}$$

At Leverage of 20, a 5% Fall Wipes Out the Equity

A bank that has lost its equity cannot lend. Selling assets to rebuild it pushes their price down for everyone else.

Liquidity Versus Solvency

Illiquid
$A_t > L_t$ at fundamental value, and the bank still cannot pay. The assets are worth more than the liabilities but cannot be sold today.

Insolvent
$A_t < L_t$ at fundamental value, even held to maturity. No amount of lending fixes that; somebody takes the loss.

Lend to the Illiquid, Resolve the Insolvent
Lending to an insolvent bank only enlarges the eventual loss. A fire sale turns an illiquid bank into an insolvent one.

Handing Over to Section 2

Section 2 takes over, with households and firms doing the choosing. Four things the model has assumed and never explained:

1. **Expectations.**
   Every result since 1.2 turns on π_t^e, unexplained.
2. **Investment.**
   $\bar{I}$ has worn a bar since the Keynesian cross.
3. **Consumption.**
   $C = c_0 + c_1(Y - T)$ has no household choosing.
4. **Price Setting.**
   The Phillips curve slope γ was never derived.

Why the Financial Sector Is in the Model

Three results, each of which survives outside the model:

1. **Why Banks Are Fragile.**
   Depositors and projects want different horizons, and pooling them is worth 28% to the impatient.
2. **Runs.**
   $r_1 > 1$ is both what makes the insurance worth having and what makes the queue worth joining.
3. **The Financial Cycle.**
   Credit and leverage peak before output does. Illiquidity and insolvency need opposite cures.

Before Lecture 1.6

1. **Play:**
   The Diamond–Dybvig app. Raise r_1 above the optimum and watch the run condition bite sooner.
2. **Apply:**
   No new problem set. Revise for the week-6 midterm, with sets 1 to 3 as the guide.
3. **Review:**
   Romer, chapter 10, on Brightspace, alongside Diamond's exposition.

> **Muddiest Point**
> One line before you go: what was the muddiest point today? The answers open Lecture 1.6.

Thank you

sam.deegan@ucdconnect.ie
sam-deegan.com

BACKUP

Appendix

Notes and Tools

- **Whelan's lecture notes** — the notes for this module. Diamond's exposition is on Brightspace.
- **The toy models** — the Diamond–Dybvig app reproduces the sample paper's numbers, nothing to install.
- **The BIS credit statistics** — credit to the private non-financial sector, the source for the credit-to-GDP exhibit.

> **Diamond Is the Only Required Reading**
> Diamond's exposition is the only required reading for this lecture.

References

Diamond, D. W. (2007). Banks and liquidity creation: A simple exposition of the Diamond–Dybvig model. *Federal Reserve Bank of Richmond Economic Quarterly*, *93*(2), 189–200. https://www.richmondfed.org/publications/research/economic_quarterly/2007/spring/diamond

Diamond, D. W., & Dybvig, P. H. (1983). Bank runs, deposit insurance, and liquidity. *Journal of Political Economy*, *91*(3), 401–419. https://doi.org/10.1086/261155