LECTURE 2.1

Rational Expectations and Investment

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

Section 2 Derives What Section 1 Assumed

Section 1 is behind us. Section 2 goes back and derives the three equations Section 1 assumed.

Last Week's Lecture

Before the midterm we closed Section 1 with banks: why they are fragile, and what a run does.

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.

Lecture Plan

1. **Rational Expectations.**
   Pin down what it means for expectations to be model-consistent.

2. **Asset Prices.**
   Derive what an asset is worth when expectations are rational.

3. **Tobin's q.**
   Turn the value of a firm into a theory of its investment.

> **Whelan, Chapters 5 and 9**
> Reading: Whelan (2023, chs. 5 and 9).

SECTION 1

Rational Expectations

What We Have Assumed So Far

This section pins down what it means for expectations to be model-consistent.

Section 1 built a model of the short run out of three equations. The Phillips curve carried a term the model never explained, inflation expectations:

$$\pi_t = \pi_t^e + \gamma(y_t - y_t^*) + \varepsilon_t^\pi \tag{1}$$

Every result in Section 1 turned on π_t^e. Shock persistence, the value of credibility, the 1970s against 2022 — all of it. Section 1 never said where π_t^e comes from.

> **Optimisation Needs a Forecast First**
> Section 1 took the three equations as given, and Section 2 derives them. You cannot derive behaviour from optimising agents without saying what they expect.

The Rational Expectations Assumption

People's forecasts are the model's own, formed from what they know at t:

$$x_{t+1}^e = \underbrace{E_t[x_{t+1}]}_{\text{the conditional expectation}} \equiv E\left[x_{t+1} \mid \Omega_t\right] \qquad (2)$$

- x_{t+1}^e: what people expect any variable x_{t+1} to be
- Ω_t: the information set, everything known at t
- $E_t[\cdot]$: the expectation formed from Ω_t

> **The Subscript on E Is a Date, Not a Person**
> E_t is not agent t's expectation. The operator uses everything known at time t. So E_{t+1} knows strictly more than E_t.

People Use the Model, Not a Rule of Thumb

Rational expectations is the claim that the two sides of that equation are the same object. Three things follow immediately:

- **Forecasts Are Endogenous:** They come out of the model rather than being fed into it, so a change in policy changes them.
- **There Is One Forecast, Not One per Person:** Everyone conditions on the same Ω_t, which is what lets us write a single π_t^e.
- **Revision Follows News, Not Events:** A forecast changes the day the news arrives.

> **The Only Internally Consistent Forecast**
>
> This is why Section 2 can derive behaviour at all. An agent optimising over the future needs a forecast of it, and rational expectations is the only forecast that does not contradict the model the agent is standing in.

Rational Expectations Is a Claim About the Error

Rational expectations does not say people are never wrong. It says their mistakes are unforecastable:

$$\eta_{t+1} = \underbrace{x_{t+1} - E_t[x_{t+1}]}_{\text{the forecast error}}, \qquad E_t[\eta_{t+1}] = 0 \tag{3}$$

- η_{t+1}: the gap between what happens and what was expected

What It Does and Does Not Claim

1. **Not That People Are Right.**
   The forecast error is almost never zero, and it can be very large.
2. **Not That People Know the Model.**
   Behaviour looks as though they did, which is a claim about outcomes rather than about heads.
3. **That Errors Are Unpredictable.**
   The error averages zero and is uncorrelated with anything known at t.

Adaptive Expectations Are the Thing This Rules Out

Adaptive expectations set this period's forecast to last period's outcome:

$$\pi_t^e = \pi_{t-1} \tag{4}$$

Wrong in the Same Direction, Every Period

Through a disinflation, inflation is falling, so π_{t-1} is above π_t every single period. The rule over-predicts every time. That is a systematic, forecastable error, and it is exactly what $E_t[\eta_{t+1}] = 0$ forbids.

Friedman's Warning Is Now an Assumption

This is Friedman (1968)'s argument from Lecture 1.2, written as an assumption rather than as a warning. There it was a prediction made in 1968 and borne out by the 1970s; here it is the only internally consistent option.

Which Errors Does Rational Expectations Allow?

Two minutes on your own, without looking, then compare with the person beside you.

> **Question**
> Which forecast errors does rational expectations allow, and which does it forbid?

Take answers from the room, and settle disagreements against $E_t[\eta_{t+1}] = 0$.

SECTION 2

Asset Prices

Why Investment, and Why Now

This section derives what an asset is worth when expectations are rational.

Investment is the smallest of the three big components of demand and much the most volatile. In a recession it does most of the falling, which is why it gets a lecture of its own and consumption gets the next one.

It is also the most forward-looking decision a firm makes. A machine bought today earns its money back over a decade. Nothing about that decision is settled by what is happening now.

> **Forward-Looking Is Why Investment Is Volatile**
> Investment is volatile **because** it is forward-looking: it rises and falls with the forecast, and forecasts shift far in a quarter.

Investment Falls Furthest and Comes Back Slowest

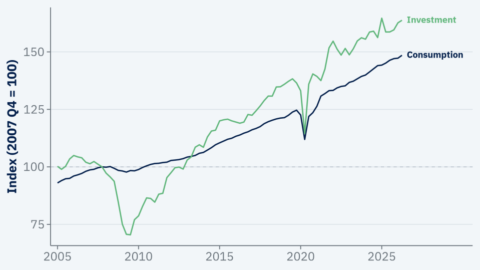

Figure 1: Consumption and investment through the last two recessions, indexed.

> **Investment Falls Further, Recovers Slower**
>
> **How to read it.** Both series are indexed to their pre-recession peak, so the comparison is of proportional falls rather than of levels. Consumption dips; investment collapses and takes far longer to come back.

The Value of a Firm

A share must earn the return available elsewhere (r), or nobody holds it:

$$P_t = \frac{1}{1+r} E_t\big[\underbrace{D_{t+1} + P_{t+1}}_{\text{the one-period payoff}}\big] \tag{5}$$

- P_t: the price of one share at t, so P_{t+1} is its resale price
- D_{t+1}: the dividend the share pays at $t+1$

> **Both Parts of the Payoff Are Unknown at t**
> Neither the dividend nor the resale price is known at t. The whole bracket therefore sits inside E_t, not just part of it.

Today's Price Is a Forecast of Tomorrow's

The same relation holds at $t + 1$, with P_{t+1} on the left:

$$P_t = \frac{1}{1+r} E_t[D_{t+1} + P_{t+1}] \qquad P_{t+1} = \frac{1}{1+r} E_{t+1}[D_{t+2} + P_{t+2}] \tag{6}$$

- The price equation is a **first-order stochastic difference equation**, the object tutorial 1 is about.
- It carries **no fundamentals yet**: the argument is pure no-arbitrage.
- It closes only once the discounted expected price vanishes in the **far future**.

> ***r* Is Held Constant**
> A time-varying required return is realistic, changes none of the argument, and costs a subscript on every term. We keep r fixed.

The Dividend Discount Model

Put the price equation at $t+1$ where P_{t+1} stood:

$$P_t = \frac{1}{1+r} E_t \left[D_{t+1} + \frac{E_{t+1}[D_{t+2} + P_{t+2}]}{1+r} \right] \tag{7}$$

Do that n times, with k counting periods ahead, for any finite n:

$$P_t = \sum_{k=1}^{n} \frac{E_t[D_{t+k}]}{(1+r)^k} + \underbrace{(1+r)^{-n} E_t[P_{t+n}]}_{\text{the terminal term}} \tag{8}$$

Why the Substitution Is Legal

The law of iterated expectations collapses $E_t\left[E_{t+1}[x_{t+2}]\right]$ to $E_t[x_{t+2}]$: today's best guess of tomorrow's best guess is today's best guess.

The Fundamental Value Is What Survives the Limit

Now let $n \to \infty$. Everything turns on whether the discounted expected price vanishes:

$$\lim_{n\to\infty} (1+r)^{-n} E_t[P_{t+n}] = 0 \quad \Longrightarrow \quad P_t = \sum_{k=1}^{\infty} \frac{E_t[D_{t+k}]}{(1+r)^k} \tag{9}$$

Killing the Last Term Is an Assumption

If that limit is not zero, the share is worth more than its dividends will ever pay: a rational bubble. Imposing zero leaves the **fundamental value**.

A Share Is Just Discounted Dividends

A share is worth its expected dividends, discounted. The rest of this lecture applies that sentence to a firm's capital rather than to its shares.

Gordon Growth

Assume dividends are expected to grow at a constant rate (g) from the dividend just paid. Then $E_t[D_{t+k}] = D_t(1+g)^k$ and the sum is geometric:

$$P_t = \frac{E_t[D_{t+1}]}{\underbrace{(r-g)}_{\text{the growth-adjusted discount rate}}} = \frac{D_t\,(1+g)}{r-g} \tag{10}$$

The Formula Needs $r > g$
If dividends grow faster than they are discounted, the sum diverges and the firm is worth infinitely much. No growth rate stays above r for ever.

What Gordon Growth Says About Prices

An entire firm, priced with two numbers and a growth rate. Only three things can raise or cut the price, and each is a claim about the future:

1. **The dividend.**
   A rise in D_t that is expected to last raises the price one for one.

2. **The growth rate.**
   A rise in g raises the price twice over: the numerator grows and $r - g$ shrinks at the same time.

3. **The required return.**
   A rise in r cuts the price, and cuts it hardest where $r - g$ was already small.

Long-Duration Assets Fall Furthest When Rates Rise

Higher Rates Hurt the Growth Firm Most

A growth firm holds most of its value in dividends twenty years away; a mature firm pays the same dividend for ever. When the required return rises, the growth firm loses far more: its distant dividends carry the largest powers of $(1 + r)$.

An Unprofitable Firm Can Still Be Valuable

Nothing here refers to current profits. A share price rises and falls with the forecast. A firm making no money can be worth a great deal.

What Each Symbol in Gordon Growth Claims

Two minutes on your own, then compare with the person beside you.

> **Question**
> Write the Gordon growth formula from memory, then say what each symbol claims.

Take answers from the room, one symbol at a time.

SECTION 3

Tobin's q

Investment and the Value of Capital

This section turns the value of a firm into a theory of its investment.

A firm should buy a machine when the machine is worth more than it costs. What the machine is worth depends on profits that have not happened yet. Tobin's answer, in 1969: the stock market has already priced them.

1. **A firm is a bundle of capital.**
   Pricing the firm therefore prices the capital inside it.

2. **The share price is a rational expectation.**
   Everything the previous section derived is already in the share price.

3. **So the decision needs one input.**
   What the market says installed capital is worth, against what new capital costs.

The Market Has Already Done the Forecasting

An Airline Worth Twice Its Fleet Should Buy

An airline whose shares are worth twice the replacement value of its fleet should be ordering aircraft. One whose shares are worth half should be letting its fleet age and not replacing it. No forecast of passenger numbers is needed by anyone except the people setting the share price.

Managers Read a Ratio, Not the Future

That is the whole of Tobin's idea, and the rest of this section is arithmetic. The manager does not have to out-forecast the market; the manager has to read one ratio off it and act.

The q Model

Tobin's average q (q_t^a) is the ratio the manager reads:

$$q_t^a = \frac{V_t}{\underbrace{p_t^K K_t}_{\text{the replacement cost of the capital stock}}} \tag{11}$$

- V_t: the market value of the firm, equity plus debt
- K_t: the capital stock in units, each new unit priced p_t^K

> **Replacement Cost, Not Historic Cost**
> The denominator is what the capital would cost today, not what the firm paid. A balance sheet gives the second; q needs the first.

A q Above One Says Buy, Below One Says Wear It Out

Read the ratio against one, and every case is a statement about what the firm should do next:

- $q_t^a > 1$: the market values installed capital above what a new set would cost, so buying a unit adds more to the firm than it takes out.
- $q_t^a < 1$: the firm should let its capital depreciate away and not replace it.
- $q_t^a = 1$: replace what wears out, and nothing more.

> **That Reading Is Average q, Not Marginal**
> The ratio just defined is average q; the decision is about the *next* unit, which is marginal q. The next two slides separate the two.

Marginal q and Average q Are Not the Same Number

The theory prices the next unit; the data price the whole firm:

$$q_t^m = \frac{1}{p_t^K}\frac{\partial V_t}{\partial K_t}, \qquad q_t^a = \frac{V_t}{p_t^K K_t} \tag{12}$$

- q_t^m: marginal q, what the next unit of capital is worth per euro it costs

A Superscript Here Names a Concept, Not a Date

Both are divided by p_t^K, so each is a pure number and the two can be compared: $\partial V_t/\partial K_t$ on its own is a price, not a ratio.

The Theory Wants q_t^m, We Have q_t^a

Every statement the theory makes is about q_t^m. Every number anyone has ever computed is q_t^a.

When Average q Stands In for Marginal q

Hayashi (1982) gives three conditions under which $q_t^m = q_t^a$ exactly:

- the firm is a **price taker** in its output and input markets;
- the **production function** has constant returns to scale;
- the **adjustment cost function** has constant returns to scale.

> **Every q Regression Uses the Wrong One**
> Marginal q cannot be measured, so the literature uses average q. The Hayashi conditions license that, and they fail wherever market power capitalises rents.

Adjustment Costs, and Why Investment Is Gradual

Installing capital uses up resources beyond the purchase price. Write those costs as quadratic in the rate of investment, ϕ the cost of installing quickly:

$$\Phi_t = \frac{\phi}{2}\left(\frac{I_t}{K_t}\right)^2 K_t \tag{13}$$

- I_t: net investment during period t
- Φ_t: resources used up installing it

Convex in the Rate, Not the Level
Doubling the capital stock over ten years is cheap; doubling it this afternoon is not.

q Pins Down the Investment Rate, Not the Capital Stock

Maximising the firm's value against the adjustment cost gives the investment rule, in marginal q — which Hayashi lets average q stand in for:

$$\frac{I_t}{K_t} = \frac{\overbrace{(q_t^m - 1)}^{\text{the excess of value over cost}}}{\phi} \qquad \text{Hayashi: } q_t^m = q_t^a \tag{14}$$

New Capital Is the Numeraire

Keep the price of new capital and the rule reads $I_t/K_t = p_t^K\,(q_t^m - 1)/\phi$ — one more symbol, the same argument.

ϕ Sets How Fast the Firm Adjusts

As $\phi \to 0$ the firm jumps straight to its desired capital stock. As ϕ grows, the investment rate falls towards zero.

The Investment Rule Is Already a Regression

Observable average q goes in where unobservable marginal q was:

$$\frac{I_t}{K_t} = \frac{\boxed{q_t^a} - 1}{\phi} \tag{15}$$

Relabel the constant and the slope, and add a residual:

$$\frac{I_t}{K_t} = a + b\, q_t^a + u_t, \qquad a = -\frac{1}{\phi}, \qquad b = \frac{1}{\phi} \tag{16}$$

- a, b, u_t: the intercept, the slope and the residual

> **Two Restrictions, Not Just a Sign**
> ϕ is free, so the theory fixes neither a nor b on its own. What it does fix: $b > 0$, and $a = -b$.

What the Regression Finds

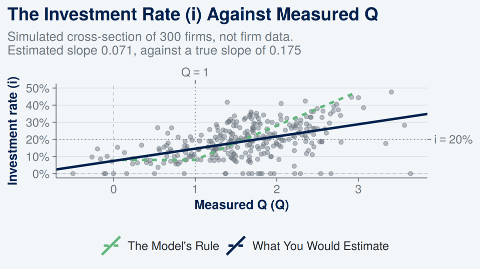

Figure 2: The model's own cross-section: investment rates against measured q for 300 simulated firms, with the investment rule that generated them (green dashed) against the fitted regression line (navy solid). The fitted line is visibly the flatter of the two, and that gap is the point.

> **The Sign Is Right and the Fit Is Poor**
> Every point here is a simulated firm, not a firm-year of data. The estimated slope comes back at 0.071 against a true slope of 0.175 — the same attenuation the real regressions report.

Why the q Regression Is Hard

Start with the regressor. Both problems below attenuate the slope — they push b towards zero whatever ϕ really is:

1. **Average q is not marginal q.**
   The regressor is a proxy, and the Hayashi conditions that justify it fail for firms with market power.

2. **Measurement error.**
   The denominator needs the replacement cost of the capital stock, which is reconstructed from accounting data. Error in a regressor biases its coefficient towards zero.

The Two Biases Pull in Opposite Directions

Perfect measurement would not rescue the regression, and the second bias inflates b:

1. **Prices that are not fundamentals.**
   Share prices that depart from the dividend discount value put noise in q_t^a — attenuation again.

2. **Simultaneity.**
   Investment and market value are set together by the same news, so both rise whatever ϕ is. That **inflates** the slope.

> **The Two Biases Cannot Be Separated**
> A small b implies a huge ϕ **or** a badly measured q, and the coefficient alone cannot say which. That is why q theory survives.

Cash Flow and Investment

Under q theory q_t^m is a sufficient statistic: nothing else about the firm should help predict investment. Add cash flow and its coefficient is significant:

- **Fazzari, Hubbard and Petersen (1988):** The cash-flow term is far larger for firms their dividend policy marks as constrained.
- **The Mechanism:** External finance costs more than internal finance, because a lender cannot see what a manager can.
- **The Objection:** Cash flow may forecast future profits better than a mismeasured q_t^a does, which is an identification problem.

An Oil Shock Breaks the Tie

A Clean Test of Financing Frictions

The 1986 collapse in the oil price cut the cash flow of oil companies. Lamont looks at the **non-oil** subsidiaries those companies owned, whose own investment prospects had not changed at all. They cut investment too.

Financing Frictions Are Real, Not Noise

The shock cut internal funds and left opportunities untouched. No better forecast explains the cut in investment, so q_t^m is not a sufficient statistic.

Investment Arrives in Bursts, Not in a Trickle

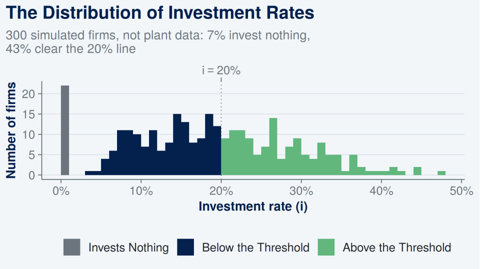

Figure 3: The model's own cross-section: investment rates across the same 300 simulated firms, the zero bar in grey and every bar above the 20 per cent spike threshold in its own colour. The empirical counterpart is the plant-level evidence this frame cites, which is not what is drawn here.

> **The Smooth Rule Predicts the Wrong Shape**
>
> The mass near zero is inaction; the long right tail is investment spikes. The smooth rule predicts a tight distribution around one point instead.

The Quadratic Cost Is the Assumption the Data Reject

Differentiate the quadratic cost and evaluate it at zero investment:

$$\left.\frac{\partial \Phi_t}{\partial I_t}\right|_{I_t=0} = 0 \tag{17}$$

The marginal cost of the *first* euro of investment is zero, so a firm with $q_t^m \neq 1$ always wants to do at least a little. Inaction is never optimal under a quadratic cost.

The Shape of the Cost Is the Problem

That property is exactly what makes the rule smooth, and it is exactly the property the plant-level data reject. The problem is not adjustment costs. It is this **shape** of adjustment cost.

What Lumpiness Changes

Installing capital stops a production line, retrains workers and extends a building. Those costs are partly **fixed**, and a fixed cost makes inaction optimal.

1. **The firm waits, then jumps.**
   The firm sits still until the gap to its desired capital is worth the fixed cost, then closes that gap at once.

2. **The response to q is state-dependent.**
   Two firms with the same q_t^m invest differently depending on how far each sits from its trigger, so no single ϕ fits both.

3. **The aggregate depends on the distribution.**
   After a long lull many firms sit bunched near their triggers, so a shock produces a burst the smooth model misses.

Lumpy Firms Need Not Make a Lumpy Aggregate

With enough heterogeneity across firms, and triggers spread widely enough, aggregate investment can behave close to the smooth benchmark even though no individual firm does.

> **Wrong for Firms, Maybe Fine in Aggregate**
> The smooth rule is wrong about firms and may be adequate about aggregates. Whether lumpiness matters at business-cycle frequencies is still argued.

Expectations, and What They Buy Us

One assumption, and the theory of investment that follows from it:

1. **Rational Expectations.**
   Forecasts are model-consistent, which lets us solve forward to a finite answer.

2. **Asset Prices.**
   A firm is worth its expected dividends, discounted, and Gordon is that value in closed form.

3. **Tobin's q.**
   Invest when the market values installed capital above what a new unit costs.

Before Lecture 2.2

1. **Play:**
   The Tobin's Q app. Turn the lumpy-firms control on and off, and watch aggregate investment after a shock.
2. **Apply:**
   Problem set 4, due before Lecture 2.2. It covers today.
3. **Review:**
   Whelan, chapter 6, on consumption.

Next Week

Today derived investment from a firm that optimises.

Consumption Generalises the 1.1 Function

Lecture 2.2 derives consumption from a household that optimises. The Keynesian consumption function from Lecture 1.1 is then a special case.

Thank you

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

BACKUP

Appendix

Notes and Tools

- **Whelan's lecture notes** — the notes for this module. Chapters 5 and 9 go with today.
- **The toy models** — Tobin's Q, with the lumpy-firms control, in your browser, nothing to install.

> **Whelan Is the Only Required Reading**
> Whelan is the only required reading. Tutorial 1 covers first-order stochastic difference equations, which is the technique behind the dividend discount model.

References

Friedman, M. (1968). The role of monetary policy. *American Economic Review*, *58*(1), 1–17.

Whelan, K. (2023). *Lecture notes on macroeconomics*. University College Dublin. https://www.karlwhelan.com/Macro2/Whelan-Lecture-Notes.pdf