LECTURE 1.1

Time Series and Macroeconomics

ECON42240 Advanced Macroeconomics

Sam Deegan
University College Dublin

Spring 2027

ECON42240 Advanced Macroeconomics

- A model earns its keep by **reproducing evidence it was never fitted to**.
- The trend, the shock and the gap are **model outputs**, not measurements.
- What a model leaves out is **where the next crisis arrives**.

What This Module Builds

Most lectures take one stage of building, estimating and judging one model class:

$$\text{model} \rightarrow \text{linearise} \rightarrow \text{solve} \rightarrow \text{state space} \rightarrow \text{likelihood} \rightarrow \text{estimate} \rightarrow \text{use}$$

Today Builds the Evidence a Model Must Match

Lectures 1.2 and 1.3 estimate the economy's responses to shocks. Lectures 2.4 and 2.5 test whether an estimated model reproduces those responses.

Where We Are

1. Time Series Methods

1.1 Time Series and Macro

1.2 Vector Autoregressions

1.3 VARs: Applications

1.4 The Kalman Filter

1.5 Solving RE Models

2. Structural Models

2.1 The RBC Model

2.2 The Phillips Curve

2.3 The New Keynesian Model

2.4 Estimating DSGE Models

2.5 The Smets-Wouters Model

3. Finance and Policy

3.1 Credit and Banking Crises

3.2 Fiscal Policy, and What Next

Lecture Plan

1. **Syllabus.**
   What the module covers, and how you are assessed.

2. **Trends and Cycles.**
   Why a trend is a modelling choice.

3. **Shocks and Propagation.**
   Where a series gets its volatility.

4. **From One Series to Many.**
   What a vector autoregression adds.

SECTION 1

Syllabus

Learning Outcomes

By the end of the module you should be able to:

- **Identify Shocks:** Estimate a VAR, defend its identification, read the responses.
- **Recover the Unobserved:** Put a model in state-space form, extract potential output.
- **Solve a Model:** Derive the conditions, log-linearise, solve the policy function.
- **Judge a Model:** Say what a mismatch with the VAR evidence means.

The final examines two more outcomes. You explain why each model replaced the last, and appraise a paper.

Methods Block Timetable

Five methods lectures come first, then the class test:

Week	No.	Lecture	Set due	Weight
1	1.1	Time Series and Macro	—	—
2	1.2	Vector Autoregressions	—	—
3	1.3	VARs: Applications	✓	5%
4	1.4	The Kalman Filter	—	—
5	1.5	Solving RE Models	✓	5%
6	—	**Class test**	—	**30%**

Models and Policy Timetable

Seven lectures on models and policy follow, then the final:

Week	No.	Lecture	Set due	Weight
6	2.1	The RBC Model	—	—
7	2.2	The Phillips Curve	✓	5%
8	2.3	The New Keynesian Model	—	—
9	2.4	Estimating DSGE Models	✓	5%
10	2.5	The Smets–Wouters Model	—	—
11	3.1	Credit and Banking Crises	✓	5%
12	3.2	Fiscal Policy, and What Next	✓	5%
—	—	**Final exam**	—	**45%**

Assessment

Three components make up the grade:

- **Problem Sets:** Six, 5% each, best five count, due two weeks after issue.
- **Class Test:** Week 6, one hour, in person, on paper, closed book. 30%.
- **Final:** Two hours, in person, on paper, **open book**. 45%.

Best Five of Six Replaces Extensions
The dropped set absorbs one missed deadline. Sets go in through the online portal, timestamped, with a sharp cut-off.

Estimated Workload

Most of the module's hours are your own time:

24

hours of lectures, two hours a week

130

hours of your own, about nine a week in term

In each problem set you estimate a model in R and judge it. The exams ask for the judgement without the computer.

Materials

This module has no textbook, so the reading is set block by block:

- **Lectures 1.1 to 1.5:** Whelan's ECON41620 notes, and Cochrane (Cochrane, 2005).
- **Lectures 2.1 to 3.2:** Romer for the models, Galí for the New Keynesian block.

Essential articles are on Brightspace and in each deck's references. All code is in R.

Using AI

I expect you will use AI on the problem sets. The rules differ by component:

- **Problem Sets:** Code assistance is permitted. Declare what you used and what for.
- **Class Test and Final:** In person, on paper. No AI, by construction.
- **What AI Cannot Do:** Judge whether an estimate is credible, which the exams ask.

How This Module Differs from the Autumn

The autumn module solved its models; this one also estimates them:

- **Front-Loaded:** Harder than the autumn module, and steepest in Lectures 1.1 to 1.5.
- **You Write the Code:** Every fortnight, in R, not watching me do it.
- **Ask Early:** After the lecture or at office hours, from week two on.

Acknowledgements

These slides build on **Karl Whelan's** ECON41620 notes, set for Lectures 1.1 to 1.5. The module keeps their shape: methods first, then the models, then finance and policy. Any errors that remain are mine.

SECTION 2

Trends and Cycles

Ireland’s 2015 Output Jump: Trend or Cycle?

Measured Irish GDP rose about 25 per cent in 2015.

> **Question**
> Did Ireland’s trend output rise 25 per cent that year? How would you know?

Splitting a Series into Trend and Cycle

Macroeconomists split a series in two, and business cycle analysis studies the cycle:

The Trend Is Non-Stationary
Where output would sit undisturbed. It has no level it returns to.

The Cycle Is Stationary
The deviation from the trend. It fluctuates around zero without drifting.

Neither part is observed, so both come out of a model.

The Log-Linear Trend

The simplest split fits a straight line with growth rate g to log output:

$$y_t = \underbrace{\alpha + gt}_{\text{the trend}} + \epsilon_t \tag{1}$$

where:

- y_t: log real output
- α: intercept
- t: period counter
- ϵ_t: stationary cyclical component

Lower Case Is a Log

Lower case marks a log throughout the module. A small deviation then reads as a percentage.

US Output Against a Log-Linear Trend

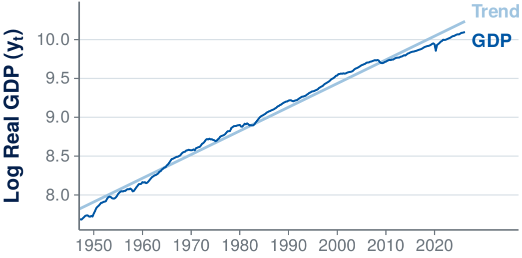

Figure 1: Log real GDP and a fitted linear trend, United States, 1947–2026.

> **The Cycle Is Small Beside the Trend**
> Log GDP rises by 2.4 over the sample. The line never misses it by more than 0.2.

The Straight-Line Residual Swings for Decades

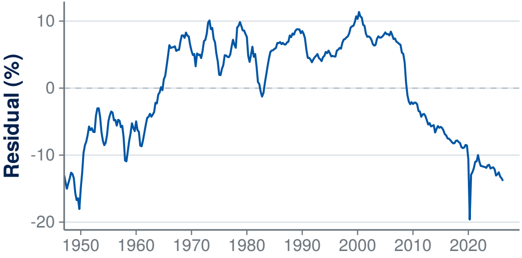

Figure 2: Residual from the fitted linear trend, United States, 1947–2026.

> **Too Long to Be a Business Cycle**
> The residual stays on one side of zero for 17 to 26 years at a time.

A Random Walk Has No Trend to Return To

Give last period's level a coefficient of one, a unit root, and add drift. Subtracting y_{t-1} from both sides gives growth:

$$y_t = g + y_{t-1} + \epsilon_t \qquad \Longrightarrow \qquad \Delta y_t = g + \epsilon_t \tag{2}$$

where:

- Δy_t: growth rate

> **A Trend Is a Model, Not a Measurement**
> Every shock is permanent, so the level never reverts to a line. A fitted line still leaves a residual that may look cyclical.

The Hodrick–Prescott Filter

Growth rates vary across decades, so the trend (y_t^*) minimises two costs over T periods:

$$\min_{\{y_t^*\}_{t=-1}^T} \left\{ \sum_{t=1}^T \underbrace{\left(y_t - y_t^*\right)^2}_{\text{the output gap}} + \lambda \sum_{t=1}^T \overbrace{\left(\Delta y_t^* - \Delta y_{t-1}^*\right)^2}^{\text{the change in the trend's growth rate}} \right\} \tag{3}$$

Lambda Is Set, Not Estimated

The smoothness weight $\lambda = 1600$ is standard for quarterly data (Hodrick & Prescott, 1997). A higher λ leaves more as cycle.

HP Cycles and the Recession Dates

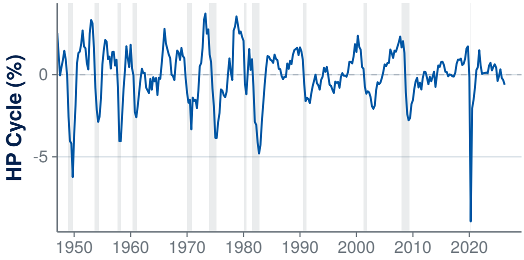

Figure 3: HP-filtered cycle in real GDP, λ = 1600, recessions shaded, 1947–2026.

> **Every Recession Since 1948 Ends in a Trough**
> The converse fails: the 2003 trough has no recession.

What the HP Filter Gets Wrong

Many DSGE modellers analyse only HP cycles. Hamilton raises three objections (Hamilton, 2018):

- **Spurious Dynamics:** Filter a random walk and a cycle appears with no source.
- **The End of the Sample:** The last few years are filtered with no data after them.
- **The Conventional λ:** Statistical criteria usually give far below 1600, leaving more as trend.

> **Policy Looks at the Least Reliable Quarters**
> Spurious dynamics in the most recent quarters make the filter least useful exactly where policymakers are looking.

Investment Swings Further Than Consumption

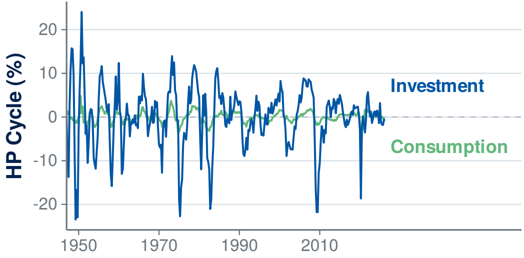

Figure 4: HP-filtered cycles in consumption and investment, United States, 1947–2026.

> **The RBC Model Is Judged on This Ratio**
> Investment's cycle has about five times consumption's standard deviation. Lecture 2.1 asks whether the RBC model reproduces that ratio.

SECTION 3

Shocks and Propagation

What Keeps a Cycle Below Trend for Several Quarters?

Five quarters in six spent below the HP trend are followed by another below it.

Question

What must a model of the cycle contain for one shock to last several quarters?

The AR(1) Model

The filter's trend is inaccurate, but its cycle is informative about growth. A first-order autoregression, AR(1), carries a fraction (ρ) of the cycle forward:

$$y_t = \rho y_{t-1} + \epsilon_t \tag{4}$$

> **A Propagation Mechanism Carries a Shock Forward**
> Here the mechanism is ρ, with $|\rho| < 1$. Today's y_t sums every past shock (ϵ_t), each faded by ρ per period.

The Two Limits of Persistence

Start the AR(1) model at $y_{t-1} = 0$, apply one shock of one, then no more:

- **At $\rho = 1$:** y stays at one for ever, as in a random walk without drift.
- **At $\rho = 0$:** The shock is gone by the next period.

RANDOM WALK MODEL

The Impulse Response Function

One unit shock from rest decays geometrically, by ρ each period:

$$y_t = 1, \quad y_{t+1} = \rho \cdot 1 = \rho, \quad y_{t+2} = \rho \cdot \rho = \rho^2, \quad \ldots \quad y_{t+h} = \rho^h \tag{5}$$

> **An Impulse Response Is an Incremental Effect**
> The extra y, h periods after one unit shock, with no later shocks.

Impulse Responses at Three Values of Persistence

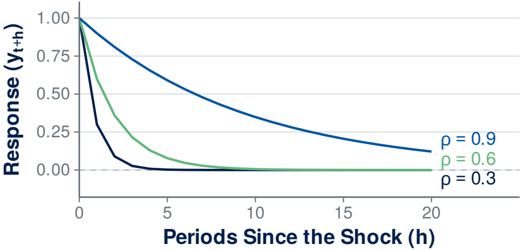

Figure 5: Response of y to a unit shock, twenty periods, $\rho = 0.3, 0.6, 0.9$.

Persistence Sets Only the Speed of Decay
Every path starts at one. At $\rho = 0.9$ about a third of the shock survives ten periods.

The AR(2) Model

Responses estimated on data do not all decay smoothly, so add a second lag:

$$y_t = \alpha + \rho_1 y_{t-1} + \rho_2 y_{t-2} + \epsilon_t \tag{6}$$

where:

- ρ_1: first-lag coefficient
- ρ_2: second-lag coefficient

> **The AR(2) Intercept Reuses the Trend's Letter**
> Here α is the AR(2) intercept. The log-linear trend used α for its own, unrelated intercept.

A Second Lag Allows Humps and Oscillations

A second lag adds two shapes to the smooth decay of an AR(1) with $\rho > 0$:

- **Hump-Shaped:** The response rises before it falls back.
- **Oscillating:** The response crosses zero and swings back over several periods.

> **Lag Length Restricts the Shapes Available**
> Practitioners choose lags with AIC, BIC and likelihood-ratio tests. Lecture 1.2 does so for a VAR.

AR(2) Shapes in the Model and in the Data

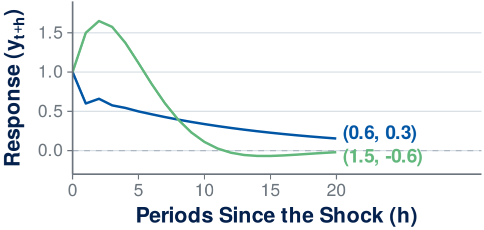

Figure 6: Model: $(\rho_1, \rho_2) = (0.6, 0.3)$ and $(1.5, -0.6)$.

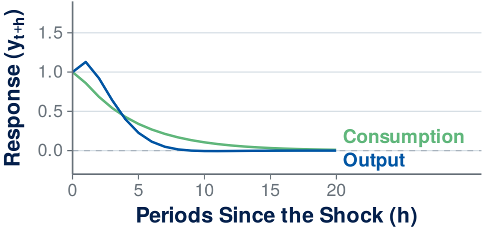

Figure 7: US HP cycles, estimated 1950–2019.

> **Output's Response Humps on Pre-Pandemic Data**
> Output's response peaks a quarter after the shock; consumption's decays at once. With 2020–26 in the sample, output's hump disappears.

Where a Series Gets Its Volatility

Volatility is the first thing measured about a cycle. Take the variance of each side of the AR(1); with $|\rho| < 1$, both dates share one variance:

$$\sigma_y^2 = \rho^2 \sigma_y^2 + \sigma_\epsilon^2 \tag{7}$$

where:

- σ_y^2: long-run variance of y
- σ_ϵ^2: shock variance

Persistence Enters Squared, with No Covariance

Scaling by ρ scales a variance by ρ^2. The shock is independent of y_{t-1}, so no covariance term appears.

Persistence Amplifies the Shock Variance

Subtract $\rho^2\sigma_y^2$, take out σ_y^2, then divide by $1-\rho^2$, positive when $|\rho|<1$:

$$\sigma_y^2 - \rho^2\sigma_y^2 = \sigma_\epsilon^2 \quad\Longrightarrow\quad \sigma_y^2\left(1-\rho^2\right) = \sigma_\epsilon^2 \quad\Longrightarrow\quad \sigma_y^2 = \frac{\sigma_\epsilon^2}{\underbrace{1-\rho^2}_{\text{the propagation term}}} \tag{8}$$

Volatility Is Shock Size and Propagation Together

The volatility of the series is partly the size of the shocks, and partly the strength of the propagation mechanism.

Shock Size and Persistence Both Raise Volatility

A larger σ_ϵ^2 raises σ_y^2, and so does a larger $|\rho|$, through $1 - \rho^2$.

> **Variance Cannot Separate Shock Size from Persistence**
> A calm economy has small shocks, or shocks that fade fast, or both. The variance alone cannot tell which.

Lectures 1.2 and 1.3 build the impulse responses that can separate the two.

The Great Moderation Began in the Mid-Eighties

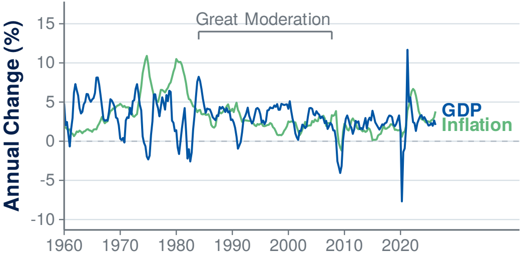

Figure 8: Four-quarter growth in real GDP and PCE inflation, United States, 1960–2026.

Both Series Calmed from the Mid-Eighties
Output and inflation became substantially less volatile after the mid-1980s. This was widely dubbed the Great Moderation.

Candidate Explanations for the Great Moderation

The candidate explanations sort under the two terms of the variance expression:

- **Smaller Shocks:** Fewer policy surprises, and smaller goods, financial and supply shocks.
- **Weaker Propagation:** Steadier policy, better inventory management, a larger services share.

> **Later Shocks Leave No Clean Period to Judge By**
> Whether 2008–09 ended the Great Moderation was already open. A pandemic and an inflation surge have landed since.

VARIANCE EXPRESSION

GDP Volatility Halved After the Mid-Eighties

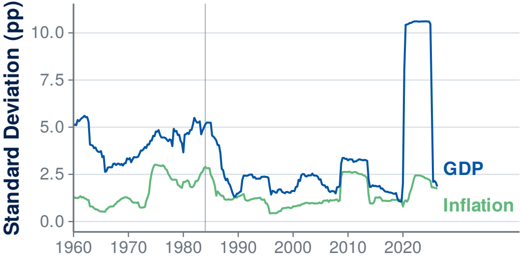

Figure 9: 20-quarter rolling standard deviations of GDP growth and PCE inflation, 1960–2026.

> **Two Pandemic Quarters Make the 2020 Plateau**
> They lift every 20-quarter window that holds them. Once they leave, GDP volatility is back near its 1990s level.

SECTION 4

From One Series to Many

Was the Volatility Since 2020 Shocks or Propagation?

GDP volatility jumped in 2020, and inflation volatility rose from 2021.

> **Question**
> Were those larger shocks, or a propagation mechanism that changed? What evidence would tell them apart?

One Series at a Time Ignores Every Other Variable

Telling shocks from propagation needs responses across many variables:

- **What an AR Misses:** Inflation and the policy rate each depend on the other's past.
- **What We Want:** The full set of responses, every variable to every shock.
- **What a VAR Adds:** Each variable depends on the lags of every variable.

The Vector Autoregression

In a vector autoregression (VAR), each of n variables depends on lags of all n (Sims, 1980). With $n = 2$ and one lag, the first equation is:

$$y_{1t} = a_{11}y_{1,t-1} + a_{12}y_{2,t-1} + e_{1t} \tag{9}$$

where:

- y_{it}: variable i
- a_{ij}: coefficient on lagged y_j
- e_{it}: VAR residual

The Cross Terms Are What Two AR(1) Models Lack

The second variable's equation uses the same two lags, with its own coefficients:

$$y_{2t} = a_{21} y_{1,t-1} + a_{22} y_{2,t-1} + e_{2t} \tag{10}$$

Two Variables Give Four Impulse Responses

The cross terms a_{12} and a_{21} carry each residual into both variables. Set both to zero and the VAR is two separate AR(1) models.

A Shock Is a Disturbance with an Economic Source

A shock is an unpredictable disturbance that starts outside the system. Four kinds are usual:

- **Policy:** Changes not explained by the systematic behaviour the VAR already captures.
- **Preferences:** Shifts in attitudes to saving or to work.
- **Technology:** Changes in the efficiency with which firms produce.
- **Frictions:** Changes in how well labour, goods or financial markets function.

Naming a Residual Does Not Identify It

A VAR residual is only what the equations failed to fit, and can mix several shocks:

> **Identification Justifies a Residual's Name**
> Calling a residual a technology or a policy shock is a claim. Justifying that claim is called identification.

Lecture 1.3 is about identification.

A VAR Describes Propagation Without Explaining It

Modern macroeconomists see fluctuations as shocks propagating through an economy. A VAR describes that propagation:

- **What a VAR Gives:** A description of how the variables move together over time.
- **What a VAR Cannot Give:** The behaviour that produced those responses.
- **Why the Gap Matters:** A description covers only policies already tried.

> **A DSGE Model Derives What a VAR Only Records**
> In a dynamic stochastic general equilibrium model, the responses come from households and firms choosing.

Before Lecture 1.2

1. **Play:**
   Pull US real GDP from FRED, detrend it with a line and with the HP filter.
2. **Apply:**
   Simulate a random walk, fit a line, and see the cycle it leaves.
3. **Read:**
   Cochrane, chapters 2 and 3, before Lecture 1.2 (Cochrane, 2005).

Next Lecture

Lecture 1.2 estimates the vector autoregression this lecture ended on:

1. **Specifying a VAR.**
   Which variables, how many lags, on what data.
2. **Estimating a VAR.**
   Equation by equation, by least squares.
3. **Reading the Output.**
   Impulse responses, error bands and variance shares.
4. **Residuals and Structural Shocks.**
   Why every mix of shocks fits the data.

Thank you

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

BACKUP

Appendix

Notation

Symbol	What it is	First met
y_t	log real output	The Log-Linear Trend
α	intercept	The Log-Linear Trend
g	trend growth rate	The Log-Linear Trend
t	period counter	The Log-Linear Trend
ϵ_t	stationary cyclical component	The Log-Linear Trend
ϵ_t	shock	A Random Walk Has No Trend to Return To
Δy_t	growth rate	A Random Walk Has No Trend to Return To
y_t^*	output trend	The Hodrick–Prescott Filter
T	sample length	The Hodrick–Prescott Filter
λ	smoothness weight	The Hodrick–Prescott Filter
y_t	the cycle	The AR(1) Model
ρ	persistence parameter	The AR(1) Model

Notation, Continued

Symbol	What it is	First met
h	periods since the shock	The Impulse Response Function
α	AR(2) intercept	The AR(2) Model
ρ_1	first-lag coefficient	The AR(2) Model
ρ_2	second-lag coefficient	The AR(2) Model
σ_y^2	long-run variance of y	Where a Series Gets Its Volatility
σ_ϵ^2	shock variance	Where a Series Gets Its Volatility
y_{it}	variable i	The Vector Autoregression
a_{ij}	coefficient on lagged y_j	The Vector Autoregression
e_{it}	VAR residual	The Vector Autoregression
n	number of variables	The Vector Autoregression
i	index of a VAR variable	The Vector Autoregression

References

Cochrane, J. H. (2005). *Time series for macroeconomics and finance*. https://www.johnhcochrane.com/s/time_series_book.pdf

Hamilton, J. D. (2018). Why you should never use the Hodrick–Prescott filter. *The Review of Economics and Statistics*, *100*(5), 831.

Hodrick, R. J., & Prescott, E. C. (1997). Postwar U.S. Business cycles: An empirical investigation. *Journal of Money, Credit and Banking*, *29*(1), 1–16. https://doi.org/10.2307/2953682

Sims, C. A. (1980). Macroeconomics and reality. *Econometrica*, *48*(1), 1–48. https://doi.org/10.2307/1912017