LECTURE 1.4

Latent Variables and the Kalman Filter

ECON42240 Advanced Macroeconomics

Sam Deegan
University College Dublin

Spring 2027

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

Last Week's Lecture

1. **The Identification Problem.**
   Why residuals are not shocks.
2. **Recursive Identification.**
   What a Cholesky ordering assumes.
3. **Monetary Policy Shocks.**
   Why the ordering can decide the answer.
4. **Long-Run Restrictions.**
   Identifying technology shocks instead.

Lecture Plan

1. **Unobserved Components.**
   Writing a model in state-space form.
2. **The Kalman Filter.**
   Extracting a state you cannot see.
3. **Estimation and Smoothing.**
   Estimating the variances, and revising the past.
4. **The Natural Rate of Interest.**
   What the filter is used for.

SECTION 1

Unobserved Components

1. Unobserved Components

2. The Kalman Filter

3. Estimation and Smoothing

4. The Natural Rate of Interest

Potential Output Is Not Measured

This section models variables no dataset contains, starting with potential output:

- **In the Model:** Inflation pressure depends on the output gap $(y_t - y_t^*)$.
- **Not in the Data:** No statistical agency measures potential output.
- **Still Restricted:** The data limit what potential output can have done.

> **A Latent Variable Is Inferred, Never Observed**
> A latent variable enters the model and never enters the data. Every estimate of one comes out of a model.

Extracting a Signal from Noise

Output jumps one quarter and inflation stays flat. Two readings compete:

- **A Large Signal:** Potential output rose, so the extra output raised no prices.
- **Mostly Noise:** Measured output and measured inflation carry quarterly errors.

> **Smooth Potential Output Makes Most of a Jump Noise**
> If potential output is smooth, most of a one-quarter jump is noise. Every method today makes this split, and each needs an assumption.

Three Estimates of the Same Output Gap

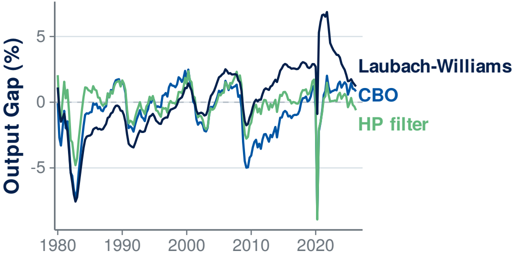

Figure 1: Output gap: CBO, HP filter and Laubach–Williams, 1980–2026, percent.

> **The Estimates Agree on Timing, Not on Size**
> In 2009Q2 the three gaps range from −1.8 to −5.0 per cent. Each comes from a different model.

Two Equations Make a State-Space Model

A state-space model holds observed and unobserved variables in two equations:

- **The State Equation:** How the unobserved variables evolve from period to period.
- **The Measurement Equation:** How the observed data depend on the unobserved.

Choosing the states, potential output among them, is the modelling decision. The modeller also supplies the two error variances.

The State Equation Carries the Unobservable

The unobserved state (S_t) evolves on its own, with a shock nobody sees:

$$\underbrace{S_t}_{\text{state}} = \underbrace{FS_{t-1}}_{\text{how it persists}} + \underbrace{u_t}_{\text{state shock}} \tag{1}$$

where:

- F: state transition matrix
- $u_t \sim N(0, \Sigma_u)$: state shock

> **Some Rows of the State Equation Are Identities**
> An identity row carries no shock, so Σ_u is not of full rank.

The Measurement Equation Links Data to States

The observed data (Z_t) are the states you lack, scaled and then mismeasured:

$$\underbrace{Z_t}_{\text{observed}} = \underbrace{HS_t}_{\text{signal}} + \underbrace{v_t}_{\text{measurement error}} \tag{2}$$

where:

- H: measurement matrix
- $v_t \sim N(0, \Sigma_v)$: measurement error

A Trend and a Cycle in State Space

Output (y_t) is a trend (y_t^*) plus a cycle, and trend growth drifts:

$$y_t = y_t^* + c_t, \qquad y_t^* = y_{t-1}^* + g_t, \qquad g_t = g_{t-1} + \eta_t \tag{3}$$

where:

- c_t: cycle
- g_t: trend growth rate
- η_t: trend growth shock

Two States and One Observable

The first equation measures output, with the cycle as its error. The states are y_t^* and g_t, so F is two by two.

The Hodrick–Prescott Filter in State-Space Form

The HP trend minimises fit plus smoothness, weighted by λ (Hodrick & Prescott, 1997):

$$\sum_t \left(y_t - y_t^*\right)^2 + \lambda \sum_t \left(\Delta y_t^* - \Delta y_{t-1}^*\right)^2 \tag{4}$$

- **The First Term:** Fit, which is the measurement equation's error c_t.
- **The Second Term:** Smoothness, which is the state equation's error η_t.

> **The HP Filter Is One State-Space Model Among Many**
> The HP filter is the previous frame's model with two variances chosen. It is a special case of this machinery.

How Much of a Surprise Belongs to the Trend?

The model is written, and the data arrive one quarter at a time.

> **Question**
> GDP comes in one per cent above the model's prediction. How much should the trend estimate absorb, and what decides it?

Take answers from the room.

SECTION 2

The Kalman Filter

1. Unobserved Components

2. The Kalman Filter

3. Estimation and Smoothing

4. The Natural Rate of Interest

Guessing a State from a Signal

The share a signal earns depends on how noisy it is. Start with one unobservable:

- **What You Want:** An estimate of a variable X that nobody observes.
- **What You Have:** An observed variable Z, correlated with X.
- **What You Need:** The two means, the covariance, and the variance of Z.

> **Joint Normality Makes the Best Guess Linear**
> X and Z are assumed jointly normal. Without normality the best guess need not be linear in Z.

The Weight You Give a Noisy Signal

Observing Z shifts your guess for X away from its mean:

$$E(X \mid Z) = \mu_X + \underbrace{\frac{\sigma_{XZ}}{\sigma_Z^2}}_{\text{the weight}} (Z - \mu_Z) \tag{5}$$

where:

- μ_X, μ_Z: unconditional means
- σ_{XZ}, σ_Z^2: covariance and signal variance

The Noisier the Signal, the Smaller Its Weight
A noisier signal gets less weight. One that covaries more with X gets more.

Many Signals Use the Same Formula

The Kalman filter applies this formula once a period, in matrices:

$$E(X \mid Z) = \mu_X + \underbrace{\Sigma_{XZ}\Sigma_{ZZ}^{-1}}_{\text{the weight matrix}} (Z - \mu_Z) \tag{6}$$

where:

- Σ_{XZ}: cross-covariance matrix
- Σ_{ZZ}: signal covariance matrix

The Filter Predicts, Then Updates

The Kalman filter works one period at a time, in two steps:

1. **Predict:** Push last period's estimate forward with the state equation alone.
2. **Update:** Revise that prediction with this period's data, then repeat.

> **Reading the Two State Estimates**
> $S_{t|t-1}$ estimates the state at t using data through $t-1$. $S_{t|t}$ is the same estimate once period t's data are in.

The Prediction Step, Before Any Data

The state equation predicts the state, then the predicted data ($Z_{t|t-1}$):

$$S_{t|t-1} = FS_{t-1|t-1} \qquad \Longrightarrow \qquad Z_{t|t-1} = HFS_{t-1|t-1} \tag{7}$$

The State Shock Has Mean Zero

u_t has mean zero, so the prediction is F times last period's estimate. $S_{t-1|t-1}$ is an estimate, not the true state.

The News Mixes a State Surprise and Measurement Error

Only the unpredictable part of period t's data is new:

$$\underbrace{Z_t - HFS_{t-1|t-1}}_{\text{the news}} \tag{8}$$

The Update Equation, Written Out

The Kalman gain (K_t) sets the share of the news the update adds:

$$S_{t|t} = \underbrace{FS_{t-1|t-1}}_{\text{prediction}} + K_t \underbrace{\left(Z_t - HFS_{t-1|t-1}\right)}_{\text{news}} \tag{9}$$

The Update Is the Conditional Expectation

$S_{t|t}$ is $E(X \mid Z)$ with $X = S_t$ and $Z = Z_t$. Under normality no other estimate has a smaller error variance.

The Kalman Gain Weighs Model Against Data

The gain sets the prediction's uncertainty ($\Sigma_{S_t|t-1}$) against the measurement noise:

$$K_t = \underbrace{\Sigma_{S_t|t-1} H'}_{\text{model uncertainty}} \left(\underbrace{\Sigma_v}_{\text{measurement noise}} + H \Sigma_{S_t|t-1} H' \right)^{-1} \tag{10}$$

The Gain Is a Ratio of Two Variances

A large Σ_v shrinks K_t, and the estimate stays near the prediction. A large $\Sigma_{S_t|t-1}$ raises K_t toward the data.

Prediction Carries Uncertainty Forward and Adds Shocks

Yesterday's uncertainty carries forward, and today's shocks add to it:

$$\Sigma_{S_t|t-1} = \underbrace{F\Sigma_{S_{t-1}|t-1}F'}_{\text{yesterday's uncertainty, carried}} + \overbrace{\Sigma_u}^{\text{new shocks}} \quad (11)$$

Prediction Need Not Raise Uncertainty
A mean-reverting F can shrink the first term below $\Sigma_{S_{t-1}|t-1}$. The shock term always adds.

Observing the Data Removes Part of the Uncertainty

Observing Z_t removes the part of the prediction variance the data explain:

$$\Sigma_{S_t|t} = (I - K_t H)\,\Sigma_{S_t|t-1} \tag{12}$$

The Gain Path Needs No Data

Neither covariance recursion uses Z_t, so the gain path is known in advance. It usually settles to a constant, the steady-state gain.

The Gain Falls as the Data Get Noisier

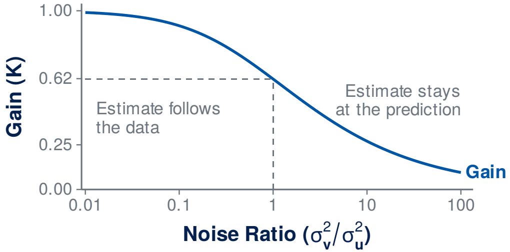

Figure 2: Steady-state gain with one state and one signal, $F = H = 1$ and $\sigma_u^2 = 1$.

> **Equal Variances Give a Gain of 0.62**
> The gain nears one as the data become exact.

Starting the Filter Off

Start a stationary state at its mean, often zero, with the covariance solving:

$$\Sigma = F\Sigma F' + \Sigma_u \tag{13}$$

> **A Random-Walk State Has No Settled Variance**
> A random-walk state, like the trend here, has no unconditional variance. Such a model starts from a very large, diffuse, initial variance.

With one state it gives Lecture 1.1's AR(1) variance.

What Happens When the Data Are Perfect

The gain figure had one state and one signal. Most models have more states than signals.

Question
Suppose the data carry no measurement error at all. Does the filter then know every state exactly?

Take answers from the room.

Perfect Data Pin Down HS_t, Not S_t

With $\Sigma_v = 0$ the update makes the filtered state fit the data exactly:

$$HS_{t|t} = Z_t \tag{14}$$

Exact Data Need Not Reveal the State
With fewer signals than states, many states fit the same Z_t. Two exact series cannot pin down three states.

With one state and $H = 1$, the estimate equals the data.

SECTION 3

Estimation and Smoothing

1. Unobserved Components

2. The Kalman Filter

3. Estimation and Smoothing

4. The Natural Rate of Interest

The Likelihood Needs States Nobody Observes

This section estimates F, H and both variances by maximum likelihood.

> **The Density of the Data Involves the Unseen States**
> The likelihood is the data's density at given parameters. Each Z_t depends on S_t, which nobody observes.

The filter's prediction stands in for S_t, with an error of known variance.

The Unseen State Enters Only Through a Forecast Error

Adding and subtracting $HS_{t|t-1}$ leaves S_t only inside a forecast error:

$$Z_t = HS_{t|t-1} + \underbrace{v_t + H\left(S_t - S_{t|t-1}\right)}_{\text{normal, with a known covariance}} \tag{15}$$

A Normal Forecast Error Gives a Likelihood

The filter computes the error's covariance from the parameters. So the data's density can be written down.

The Likelihood of One Observation

Each period's log density depends on its forecast error and that error's variance:

$$\log f\left(Z_t \mid Z_{t-1}, \theta\right) = -\frac{1}{2}\left[\log|\Omega_t| + \varepsilon_t' \Omega_t^{-1} \varepsilon_t\right] + \text{constant} \qquad (16)$$

where:

- $\varepsilon_t = Z_t - HS_{t|t-1}$: forecast error
- $\Omega_t = \Sigma_v + H\Sigma_{S_t|t-1}H'$: error covariance
- $\theta = (F, H, \Sigma_u, \Sigma_v)$: parameter vector
- $-\frac{m}{2}\log 2\pi$: the constant, for m observables

Both Terms in the Bracket Must Be Small

Inflating Ω_t shrinks the quadratic term and raises $\log|\Omega_t|$. The log density is highest where their sum is smallest.

The Likelihood of the Whole Sample

One pass of the filter gives every term, and the terms add:

$$\log f\left(Z_1, \ldots, Z_T \mid S_{1|0}, \theta\right) = -\frac{1}{2} \sum_{t=1}^{T} \left[\log|\Omega_t| + \varepsilon_t' \Omega_t^{-1} \varepsilon_t\right] + \text{constant} \tag{17}$$

The Filter Makes the Likelihood Computable

Without the filter the likelihood involves every unobserved state. An optimiser proposes θ, runs the filter, reads the sum, and repeats.

Filtering Uses the Past, Smoothing Uses Everything

Each past state has two estimates, one for each set of data:

- **The Filter:** Uses data up to period t only, as a forecaster would.
- **The Smoother:** Uses the whole sample, including data after period t.

> **One-Sided Uses the Past, Two-Sided Uses Everything**
> The filter's estimate is one-sided and the smoother's is two-sided. Researchers generally report the smoother, since the whole sample is in hand.

The Smoothed Estimate Is the Steadier One

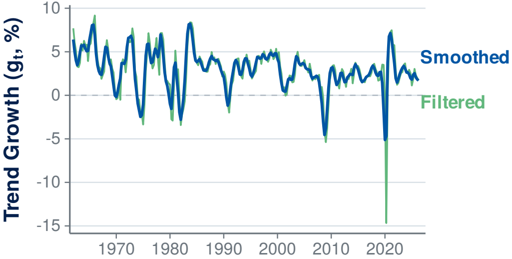

Figure 3: Filtered and smoothed trend growth (g_t) at the estimated variances, 1962–2026.

The Two Coincide at the Sample End

At 2026Q2 no later data exist, so the two agree. The smoother works backwards from there.

The HP Filter's Lambda Is a Variance Ratio

With a diffuse start, smoothing the trend-and-cycle model gives exactly the HP filter at (Hodrick & Prescott, 1997):

$$\lambda = \frac{\sigma_c^2}{\sigma_\eta^2} \tag{18}$$

where:

- σ_c^2: cycle variance
- σ_η^2: trend growth shock variance

> **Lambda Is an Assumption About Two Variances**
> Hodrick and Prescott assumed standard deviations of 5 and one eighth percentage points. So $\lambda = 5^2 \times 8^2 = 1600$.

HP IN STATE SPACE

Estimating the Variances Gives a Lambda Near 1.6

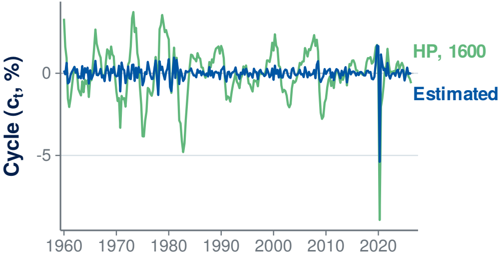

Figure 4: Cycle at the estimated variances and the HP cycle at 1600, US, 1960–2026.

> **The Estimated Trend Tracks GDP Closely**
> The likelihood peaks at a variance ratio of 1.6. That cycle is a third as volatile as HP's. Neither cycle is a measurement.

The Filter Returns in Lecture 2.4

Estimating a DSGE model runs this filter, unchanged, thousands of times:

> **A Solved DSGE Model Is a State-Space Model**
> Solving the model delivers the two equations. The filter's forecast errors then give its likelihood, the same sum as here.

Lecture 2.4 adds priors over the parameters, and a posterior.

Where Would the Natural Rate Enter a Model?

Policy is set against a real interest rate that nobody observes.

> **Question**
> A central bank wants the real rate at which output sits at potential. How would you estimate it with today's tools?

Take answers from the room.

SECTION 4

The Natural Rate of Interest

1. Unobserved Components

2. The Kalman Filter

3. Estimation and Smoothing

4. The Natural Rate of Interest

What the Natural Rate of Interest Is

This section applies the filter to the rate that policy is set against:

> **The Real Rate That Leaves Output at Potential**
> The natural rate (r_t^*) is the real rate consistent with output at potential. No market quotes it, so every estimate comes from a model.

A real rate above r_t^* restrains demand, and one below it stimulates demand.

Demand Ties the Natural Rate to the Output Gap

A real rate above the natural rate lowers next quarter's output gap (Laubach & Williams, 2003):

$$\tilde{y}_t = A_y(L)\tilde{y}_{t-1} + A_r(L)\left(r_{t-1} - r^*_{t-1}\right) + \epsilon_{1t} \qquad \textbf{(19)}$$

where:

- $\tilde{y}_t$: output gap
- r_t: real interest rate
- $A_y(L), A_r(L)$: lag polynomials

The Phillips Curve Ties the Output Gap to Inflation

A positive output gap raises inflation (π_t) the following quarter:

$$\pi_t = B_\pi(L)\pi_{t-1} + B_y(L)\tilde{y}_{t-1} + B_x(L)x_t + \epsilon_{2t} \quad (20)$$

where:

- x_t: other inflation determinants
- $B_\pi(L), B_y(L), B_x(L)$: lag polynomials

Both Are Measurement Equations
Output, inflation and the real rate are observed. The output gap and r_t^* inside them are states.

The Natural Rate Is Trend Growth Plus a Residual

Laubach and Williams scale trend growth by c and add a residual (z_t):

$$r_t^* = cg_t + z_t, \qquad z_t = D_z(L)z_{t-1} + \epsilon_{3t} \tag{21}$$

Faster Trend Growth Raises the Natural Rate

The link needs c positive, and the current estimate is 1.07. The residual collects everything else.

Potential Output Is the Local Linear Trend Again

Potential output (y_t^*) grows at a drifting rate g_t, plus a level shock:

$$y_t^* = y_{t-1}^* + g_{t-1} + \epsilon_{4t}, \qquad g_t = g_{t-1} + \epsilon_{5t} \tag{22}$$

A Separate Shock Hits the Level

Section 1's trend used one shock, η_t, for level and growth. Here ϵ_{4t} shifts the level alone.

TREND IN STATE SPACE

Two Unobservables from One Dataset

One sample has to pin down two smooth series nobody observes:

- **Potential Output:** Identified by output, and by inflation through the Phillips curve.
- **The Natural Rate:** Identified only through the demand equation's real-rate gap.

> **The Estimates Carry Wide Standard Errors**
> Laubach and Williams report wide bands around the natural rate. A point estimate alone overstates how precise it is.

The Estimated Natural Rate Fell Until 2010

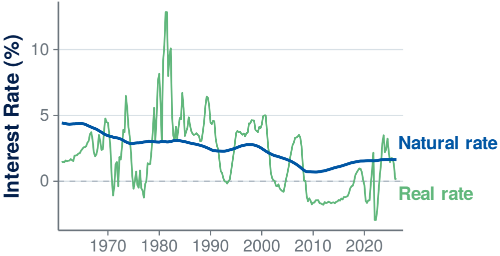

Figure 5: Laubach–Williams natural rate (two-sided) and the real policy rate, 1961–2026.

> **The Estimate Has Risen Since 2010**
> The estimate fell from 4.4 in 1961 to 0.7 in 2010. By 2026Q2 it was 1.65. New York Fed update of Laubach & Williams (2003).

Slower Trend Growth Shrinks the Estimated 2009 Gap

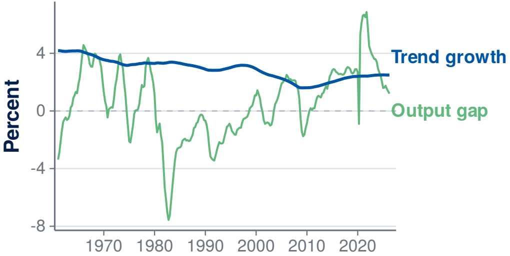

Figure 6: Laubach–Williams trend growth (g_t) and output gap, two-sided, 1961–2026, percent.

Weak Growth Read as Structural, Not Cyclical

Trend growth fell from 4.2 in 1961 to 1.6 by 2010. The 2009 gap is only −1.8. The 2020–21 values carry the model's COVID adjustment.

One-Sided and Two-Sided Estimates Differ

A policymaker in any quarter could use only data up to that quarter:

> **Only the One-Sided Estimate Was Ever Available**
> Policy in any quarter could use only a one-sided estimate. The two-sided estimate exists only once later data arrive.

Later data change the estimate of what was trend at the time.

Section 3’s filter-and-smoother frame defines both estimates.

ONE-SIDED AND TWO-SIDED

The Latest Estimate Has No Two-Sided Version

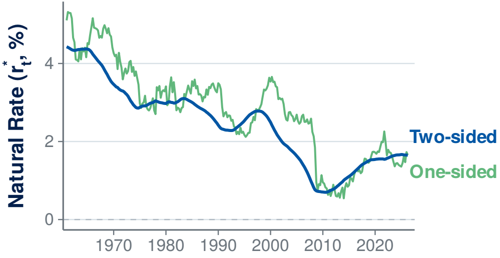

Figure 7: Laubach–Williams natural rate, one-sided and two-sided, 1961–2026, percent.

> **The One-Sided Estimate Ran 1.6 Points High in 2007**
> In 2007Q4 the one-sided estimate was 2.6 and the two-sided 1.0. At 2026Q2 the two coincide.

Real-Time Estimates Get Revised Heavily

The HP filter's latest values are its least reliable too (Hamilton, 2018):

- **The HP Version:** Recent values are filtered with no future data available.
- **The Filter Version:** A one-sided estimate uses no data after that quarter.

> **Policy Looks at the Least Reliable Quarters**
> Policy is made at the end of the sample, on one-sided estimates. Revisions arrive years after the decision.

Whether the Natural Rate Has Risen Since 2022

Estimates of the natural rate declined for decades (Holston et al., 2017). Whether real rates since 2022 signal a higher natural rate is unsettled:

- **The Case for a Rise:** Government borrowing, defence spending and capital demand.
- **The Case Against:** Demographics and productivity, which change slowly.

The two-sided estimate rose only from 1.55 to 1.65 from 2022Q1 to 2026Q2.

Which Estimate Should the Central Bank Use?

Two estimates of every past quarter's natural rate exist. Only one exists for this quarter.

> **Question**
> You advise a central bank on whether policy is currently tight. Which estimate do you put in the briefing, and what else belongs there?

Take answers from the room.

Setting Policy with a Mismeasured Natural Rate

Policy is set against an estimate, so estimation error reaches the policy rate:

- **An Estimate Too Low:** Policy is looser than intended and demand runs hot.
- **An Estimate Too High:** Policy is tighter than intended and demand is squeezed.

> **Brief with a Band and More Than One Estimate**
> The latest one-sided estimate is the least reliable. Report its band and a second model's estimate alongside it.

Before Lecture 1.5

1. **Play:**
   Pull the New York Fed's natural rate file and plot both estimates.
2. **Apply:**
   Fit a local linear trend to Irish GDP, against the HP trend.
3. **Read:**
   Whelan's Part 6 slides on rational expectations, before Lecture 1.5.

Next Lecture

Lecture 1.5 solves models in which expectations of the future determine today:

1. **Rational Expectations.**
   The assumption, and the forward solution it delivers.
2. **Reduced Forms and Persistence.**
   A solved model as a VAR, and lagged dynamics.
3. **Solving the System.**
   Eigenvalues, the Blanchard–Kahn count, and generalised Schur.
4. **The Lucas Critique.**
   Why reduced-form coefficients shift when policy changes.

Thank you

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

BACKUP

Appendix

Notation, Part I

Symbol	What it is	First met
S_t	unobserved state	The State Equation Carries the Unobservable
F	state transition matrix	The State Equation Carries the Unobservable
u_t	state shock	The State Equation Carries the Unobservable
Σ_u	state shock covariance	The State Equation Carries the Unobservable
N	normal distribution	The State Equation Carries the Unobservable
Z_t	observed data	The Measurement Equation Links Data to States
H	measurement matrix	The Measurement Equation Links Data to States
v_t	measurement error	The Measurement Equation Links Data to States
Σ_v	measurement error covariance	The Measurement Equation Links Data to States
y_t	log real output	Lecture 1.1
y_t^*	output trend, potential output	Lecture 1.1
c_t	cycle	A Trend and a Cycle in State Space
g_t	trend growth rate	A Trend and a Cycle in State Space

Notation, Part II

Symbol	What it is	First met
η_t	trend growth shock	A Trend and a Cycle in State Space
λ	smoothness weight	Lecture 1.1
Δ	first difference	Lecture 1.1
X, Z	unobserved and observed variables	Guessing a State From a Signal
$E(X \mid Z)$	best guess of X given Z	The Weight You Give a Noisy Signal
μ_X, μ_Z	unconditional means	The Weight You Give a Noisy Signal
σ_{XZ}, σ_Z^2	covariance and signal variance	The Weight You Give a Noisy Signal
Σ_{XZ}	cross-covariance matrix	Many Signals Use the Same Formula
Σ_{ZZ}	signal covariance matrix	Many Signals Use the Same Formula
$S_{t\|t-1}, S_{t\|t}$	state estimates to $t-1$ and t	The Filter Predicts, Then Updates
$Z_{t\|t-1}$	predicted data	The Prediction Step, Before Any Data
K_t	Kalman gain	The Update Equation, Written Out
$\Sigma_{S_t\|t-1}$	prediction error covariance	The Kalman Gain Weighs Model Against Data

Notation, Part III

Symbol	What it is	First met
$\Sigma_{S_t\|t}$	filtered error covariance	Observing the Data Removes Part of the Uncertainty
I	identity matrix	Observing the Data Removes Part of the Uncertainty
σ_u^2, σ_v^2	scalar state and measurement variances	The Gain Falls as the Data Get Noisier
Σ	unconditional state covariance	Starting the Filter Off
ε_t	forecast error	The Likelihood of One Observation
Ω_t	forecast error covariance	The Likelihood of One Observation
θ	parameter vector	The Likelihood of One Observation
f	density	The Likelihood of One Observation
$S_{1\|0}$	first state prediction	The Likelihood of the Whole Sample
T	sample length	The Likelihood of the Whole Sample
σ_c^2	cycle variance	The HP Filter's Lambda Is a Variance Ratio
σ_η^2	trend growth shock variance	The HP Filter's Lambda Is a Variance Ratio

Notation, Part IV

Symbol	What it is	First met
r_t^*	natural rate	What the Natural Rate of Interest Is
$\tilde{y}_t$	output gap	Demand Ties the Natural Rate to the Output Gap
r_t	real interest rate	Demand Ties the Natural Rate to the Output Gap
L	lag operator	Lecture 1.1
$A_y(L), A_r(L)$	lag polynomials	Demand Ties the Natural Rate to the Output Gap
$\epsilon_{1t}, \ldots, \epsilon_{5t}$	equation shocks	Demand Ties the Natural Rate to the Output Gap
π_t	inflation	The Phillips Curve Ties the Output Gap to Inflation
x_t	other inflation determinants	The Phillips Curve Ties the Output Gap to Inflation
$B_\pi(L), B_y(L), B_x(L)$	lag polynomials	The Phillips Curve Ties the Output Gap to Inflation
z_t	residual	The Natural Rate Is Trend Growth Plus a Residual
$D_z(L)$	lag polynomial	The Natural Rate Is Trend Growth Plus a Residual
c	loading of the natural rate on growth	The Natural Rate Is Trend Growth Plus a Residual

References

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

Holston, K., Laubach, T., & Williams, J. C. (2017). Measuring the natural rate of interest: International trends and determinants. *Journal of International Economics*, *108*, S59–S75. https://doi.org/10.1016/j.jinteco.2017.01.004

Laubach, T., & Williams, J. C. (2003). Measuring the natural rate of interest. *The Review of Economics and Statistics*, *85*(4), 1063–1070.